arXiv:quant-ph/0112149v1 23 Dec 2001

Emergence of particles from bosonic quantum field theory

David Wallace*

An examination is made of the way in which particles emerge from linear, bosonic, massive quantum field theories. Two different constructions of the one-particle subspace of such theories are given, both illustrating the importance of the interplay between the quantum-mechanical linear structure and the classical one. Some comments are made on the Newton-Wigner representation of one-particle states, and on the relationship between the approach of this paper and those of Segal, and of Haag and Ruelle.

Keywords: Quantum Field Theory; Particle Localization; Relativistic Quantum Mechanics

1 Introduction

For better or for worse, most quantum systems are found by starting with a classical system and then quantizing it. The states of the resulting quantum system will be described by complex functions on the configuration space of the classical system, whose squared moduli tell us the probability density for finding the system in a given configuration.

Applying this method to classical fields would seem, if not unproblematic, then at least difficult only for technical reasons. We would naturally expect to find a theory whose states are wave-functionals on the configuration space: that is, maps which associate a complex number to each configuration of the classical field on a given hypersurface.

Scarcely a vestige of this behaviour is seen in the usual phenomenology of 'quantum field theory'. Instead we find ourselves with a theory usually described in terms of particles: quarks, gluons, electrons..., and the localized interactions between them. The field-configuration viewpoint is occasionally seen (notably in the path-integral formalism of quantum field theory) but is usually regarded as at best a calculational tool.

(*December 23, 2001*)
* Centre for Quantum Computation, The Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, U.K.
(*e-mail:* david.wallace@merton.ox.ac.uk).

Furthermore, these particles are notoriously strange entities. Various results of quantum field theory seem strongly to imply that they cannot be localized in any meaningful, covariant way; that they must be created and annihilated in interactions which cannot be spatio-temporally localized; that we cannot start with a theory of free particles and 'turn on' an interaction without pathology (implying that the concept of particle is bound up with the dynamics of the theory and is not just a kinematic concept); and that the particles which should be associated with a given field theory vary according to the energy levels at which that theory is studied.

For these reasons, it is normal in modern quantum field theory to regard the field as the primary concept and the particles as secondary, derivative entities. This process has been studied extensively using the methods of algebraic quantum field theory and the signs are encouraging that it can be understood in a mathematically and conceptually rigorous way; however, the very abstractness of these methods can make it difficult to understand quite why the idea of 'particle' should be so powerful in understanding the prima facie very different concepts inherent in a quantised *field*.

The purpose of this paper, then, is to analyse in a fairly concrete context the way in which certain subspaces of a quantum field theory's Hilbert space come to possess characteristics of a one-quantum-particle Hilbert space. 'The concrete context' in question is that of a massive, scalar, bosonic field, assumed to be asymptotically treatable as a free field. Section 2 presents the classical and quantum theory of such a field, and section 3 considers what the correct definition should be of 'local' and 'particle' states in QFT. In sections 4–6 — which form the core of the paper — two separate constructions of the one-particle subspace are given, both of which illustrate the central role played by the interaction between the linear structure on the QFT Hilbert space (present in any quantum system) and the linear structure on the classical phase space (specific to a linear field theory). Section 7 discusses the Newton-Wigner representation of position for one-particle systems, and section 8 makes some brief comments on the relationship between this paper's approach to QFT and some other approaches; section 9 is the conclusion.

Currently, it is increasingly common for foundational discussions of QFT to be conducted in the powerful and abstract language of algebraic quantum field theory. This paper eschews that tendency: some concepts and results of algebraic QFT are referred to, but the framework used here is much closer to that used by the mainstream physics community. For a defence of the validity of using this apparently rather non-rigorous framework for a foundational discussion, see Wallace (2001b).

2 Field quantization

In this section, we shall review the method by which free, and weakly interacting, field theories are quantized. We shall outline the problems which occur when we try to reinterpret these quantized theories as fundamentally about particles, and then consider, in qualitative terms, how particles can enter the

theory in a non-fundamental way.

2.1 Classical free fields

A fairly general second-order field equation for a free-field theory is

$$\frac{\mathrm{d}^2}{\mathrm{d}t^2}\phi(\mathbf{x},t) + \mathcal{R}\phi(\mathbf{x},t) = 0 \tag{1}$$

where $\phi(\mathbf{x},t)$ is a real field defined on some manifold $\Sigma \times \mathrm{R}$, points on Σ are labelled by $\mathbf{x}$, and $\mathcal{R}$ is some real symmetric operator acting on functions of the spatial coordinate $\mathbf{x}$.

The simplest example of such a theory is the real Klein-Gordon equation, for which

$$\mathcal{R} = m^2 - \nabla^2, \tag{2}$$

and in fact if we want a theory with relativistic covariance then there are no other examples. However, more general theories of this form are potential models for:

- Curved spacetimes with a time translation symmetry;
- Systems interacting with a time-independent background field;
- Solid-state systems.

Most crucially for our purposes, field theories with this form also occur as approximations to nonlinear theories; we will consider this case in more detail later.

Such a theory can be generated from a Lagrangian

$$\mathrm{L}[\phi,\dot{\phi}] = \frac{1}{2}\int_\Sigma \mathrm{d}^3\mathbf{x}\,\left(\dot{\phi}(\mathbf{x})^2 - \phi(\mathbf{x})(\mathcal{R}\phi)(\mathbf{x})\right). \tag{3}$$

Carrying out the Legendre transform to the Hamiltonian formalism, we get a set of canonical coordinates $\phi(\mathbf{x})$ labelled by $\mathbf{x}$, a set of conjugate momenta $\pi(\mathbf{x})$, and a Hamiltonian H, where

$$\pi(\mathbf{x}) =_{df} \frac{\delta \mathrm{L}}{\delta\dot{\phi}(\mathbf{x})} = \dot{\phi}; \tag{4}$$

$$\mathrm{H}[\phi,\pi] = \frac{1}{2}\int_\Sigma \mathrm{d}^3\mathbf{x}\,\left(\pi(\mathbf{x})^2 + \phi(\mathbf{x})(\mathcal{R}\phi)(\mathbf{x})\right). \tag{5}$$

Points in the phase space $\mathcal{P}$ of the field are then given by specifying pairs of functions (ϕ,π); the Poisson bracket on $\mathcal{P}$ is given by

$$\{A[\phi,\pi], B[\phi,\pi]\} = \int_\Sigma \mathrm{d}^3\mathbf{x}\,\left(\frac{\delta A}{\delta\phi(\mathbf{x})}\frac{\delta B}{\delta\pi(\mathbf{x})} - \frac{\delta A}{\delta\pi(\mathbf{x})}\frac{\delta B}{\delta\phi(\mathbf{x})}\right), \tag{6}$$

so of course ϕ and π obey the canonical relations $\{\phi(\mathbf{x}),\phi(\mathbf{y})\} = \{\pi(\mathbf{x}),\pi(\mathbf{y})\} = 0$ and $\{\phi(\mathbf{x}),\pi(\mathbf{y})\} = \delta(\mathbf{x}-\mathbf{y})$.

Through each point in phase space flows a unique trajectory; hence points in $\mathcal{P}$ are in one-to-one correspondence with solutions of (1).

2.2 Field quantization

We will quantize classical fields (free or interacting) in the most naive possible way: by direct comparison with non-relativistic particle mechanics. That is, we will represent states of the quantum system by complex wave-functions on the configuration space of the classical system. In this case, that configuration space is the infinite-dimensional space $\mathcal{S}$ of functions on Σ, so the quantum states will be functionals $\Psi[\chi]$ on this space (we will denote the Hilbert space of all such functionals as $\mathcal{H}_\Sigma$). By analogy with the non-relativistic quantization of the coordinates q, p as

$$\widehat{q}\psi = q\psi(q); \;\; \widehat{p}\psi = -i\frac{\mathrm{d}\psi}{\mathrm{d}q} \tag{7}$$

we will quantize the coordinates $\phi(\mathbf{x})$ and $\pi(\mathbf{x})$ as

$$(\widehat{\phi(\mathbf{x})}\Psi)[\chi] = \chi(\mathbf{x})\Psi[\chi]; \tag{8}$$

$$(\widehat{\pi(\mathbf{x})}\Psi)[\chi] = -i\frac{\delta\Psi}{\delta\chi(\mathbf{x})}[\chi]. \tag{9}$$

It is easy to check that the canonical commutation relations are satisfied:

$$\left[\widehat{\phi(\mathbf{x})}, \widehat{\phi(\mathbf{y})}\right] = \left[\widehat{\pi(\mathbf{x})}, \widehat{\pi(\mathbf{y})}\right] = 0; \tag{10}$$

$$\left[\widehat{\phi(\mathbf{x})}, \widehat{\phi(\mathbf{y})}\right] = i\delta(\mathbf{x}-\mathbf{y}). \tag{11}$$

It is to be admitted that we have been very cavalier with our treatment of the infinite-dimensional spaces in use here. It is possible (whilst we confine ourselves to free fields) to be much more careful and rigorous,[1] but if we wish our framework to be powerful enough to handle interactions then there is actually no need for infinite-dimensional technicalities, for reasons to be explained in section 2.3.

2.3 Interactions and renormalisation

Formally speaking, nothing in the previous description will be altered if we add some higher-order terms (such as ϕ^4), which change the field equation from free to interacting: we could restrict our attention to regimes in which these terms are small in comparison to the free-field Hamiltonian, and proceed to analyse their effects using perturbation theory.

However, the reader may at this stage object that we are playing fast and loose with some very poorly-defined mathematical concepts. In fact, it is well-known that terms like ϕ^4, when added to the Hamiltonian, give contributions which are not small, but infinite — hence formulating a well-defined interacting

[1] See Marsden and Ratiu (1994); Woodhouse (1991) for discussions of infinite-dimensional classical mechanics, and Wald (1994) for a careful discussion of quantising linear field theories.

quantum theory is actually very subtle. In fact, one approach would be to say that the only quantum theories we understand well enough for conceptual study are the free-field ones, and confine our attention to those.

In this paper, however, we shall take a more liberal attitude. There is actually a well-defined approach to understanding these apparent infinities, worked out primarily by Kenneth Wilson and originating in solid-state physics. In Wilson's approach, we postulate that QFTs do not after all have infinitely many degrees of freedom; rather, some unknown processes cut off the high-energy degrees of freedom and leave only finitely many to contribute to the physics. It then turns out — rather remarkably — that all interaction terms in the Hamiltonian will fall into two categories. *Non-renormalisable* interactions will be negligibly weak on energy scales far lower than the cutoff threshold. *Renormalisable* interactions are not necessarily negligible, but at low energies they are affected by the choice of the cutoff only through modifications ("renormalisation") of the parameters in the interaction terms. Since these parameters are in any case only known through experiment, the choice of the cutoff becomes irrelevant to the low-energy regimes of the QFT.

Solid-state physics provides an example of this process. If we study a solid-state system on length-scales which are large compared to the interatomic spacing, we can approximate the possible (classical) configurations of the atoms by a continuous function — and thus approximate the system by a continuous field theory. In quantizing this theory we find that interaction terms lead to infinities, but these are an artefact of our continuum assumption. Once we introduce a cutoff banning excitations of the system which vary significantly on length-scales short in comparison with the interatomic separation, the infinities vanish.

Because we are understanding field theories in this way, we can take a relaxed attitude to the infinite-dimensional spaces which we will encounter in our analysis: such spaces are 'really' finite-dimensional, with the very short-distance excitations disallowed. As for the interaction terms, we will not have need of their specific forms. We shall just assume, where necessary, that such terms are present but that the theory has been renormalised and that, after renormalisation, the interaction terms can be treated perturbatively. For details of the mathematics of this process, see Peskin and Schroeder (1995) or any other QFT textbook; for a conceptual discussion see Wallace (2001b).

2.4 Problems with a particle interpretation

The theory constructed above is undeniably a field theory, in the sense that its configuration space, and fundamental observables, are inherently field-theoretic. It is, however, tempting to try to reinterpret the theory so as to make direct contact with the particle concept, either by establishing some kind of 'duality' between field and particle descriptions (in the same sense that there is a duality between position and momentum representations in ordinary quantum mechanics, with neither representation being privileged over the other) or by replacing the field description entirely with a particulate one (in which case, presumably, the field observables would just count as auxiliary constructions of

no direct physical significance).

There are however, many problems which emerge as soon as we try to interpret any QFT so as to incorporate particles at a fundamental level:

- The 'elementary particles' of particle physics are generally understood as pointlike objects, which would seem to imply the existence of position operators for such particles. However, if we add the requirement that such operators are covariant (so that, for instance, a particle localised at the origin in one Lorentz frame remains so localised in another), or the requirement that the wave-functions of the particles do not spread out faster than light, then it can be shown that no such position operators exist. (See Halvorson and Clifton (2001), and references therein, for details.)

- In non-relativistic quantum mechanics, it is straightforward to construct Hamiltonians which describe particles interacting via long-range forces (for a simple example, consider two charged particles interacting via a Coulomb force). However, the concept of a long-range interaction *prima facie* requires some sort of preferred reference frame, which seems to cast doubt upon the possibility of constructing such an interaction in a relativistically covariant way.

- As was mentioned in section 2.3, if interactions are present in a QFT then it is necessary to work, not with the bare parameters in the Hamiltonian, but with 'renormalised' parameters — and the parameters which must be renormalised include some of those, such as charge, which are generally taken to be intrinsic properties of particles. However, there is no privileged way of renormalising the parameters, so that the values of these parameters — and hence, the natures of the particles which they purport to describe — can be in part a purely conventional matter.

- When we consider quantum field theory on a general spacetime background, there is no unique procedure to define particles, and states which appear particulate in one reference frame do not do so in other reference frames. For instance, consider the so-called 'Unruh effect', in which the Minkowski vacuum of a free QFT looks like a thermal (hence, non-particulate) state to a uniformly accelerating observer. In this example it may be possible to argue that non-inertial observers' descriptions are somehow less fundamental, but in a less symmetric spacetime there will be no preferred class of observers available, hence no preferred definition of particle. (For a more detailed account of this point, see Wald (1994), who advocates abandoning the particle concept as a consequence.)

Not all have abandoned particles as fundamental in view of these difficulties: Fleming has given a strong defence of the idea that particle localisation does indeed make sense in relativistic QFT (see Fleming (1996), Fleming and Butterfield (1999), and references therein) and Weinberg's recent QFT textbook (Weinberg 1995) explicitly begins with particles and constructs the fields as auxiliary objects. However, the general consensus in QFT (insofar as such

issues are ever explicitly addressed[2]) appears to be that the subject is primarily about quantum *fields*. In fact, much modern research in the field only really makes sense from this viewpoint: for example, consider lattice quantum chromodynamics (which attempts to understand quark confinement and the existence of protons and neutrons, but is formulated in terms of field configurations and makes only limited contact with the elementary heuristic that a proton is 'just' three particulate quarks bound together); or consider the quantum Sine-Gordon equation (Coleman 1985), which has two distinct particle descriptions (one fermionic, one bosonic) with the weak-field version of the one equivalent to the strong-field version of the other).

Of course, none of this is to deny that particles *exist*, merely that they are not part of the fundamental ontology of quantum field theory. In the next section we will consider how it might be possible for the particle concept to be recovered from a field-theoretic description.

2.5 The particle as emergent concept

It is a central result of condensed-matter physics that, if we start with some macroscopic collection of nonrelativistic particles close to some collective stable state, small excitations from that state can often be treated in terms of creating 'particles'. It is also generally true that, for strongly interacting systems, these 'particles' do not coincide with the particles from which the system is built: so vibrations in a crystal are described in terms of 'phonons', which are not crystal atoms, and quantized waves in a magnet are described in terms of 'magnons' which are not iron atoms (Kittel and Fong 1987).

There are striking formal parallels with quantum field theory: in fact, the construction of phonons from a monatomic crystal is virtually the same as the construction of particle states in a massless, scalar quantum field theory. The difference is, the ontology of a crystal is not in question. It is definitely made up of the lattice atoms - which correspond to the field states at different space points in scalar QFT. Nonetheless many phenomena can be described by regarding the crystal as a gas of phonons, and some — e. g. heat transport — require us to think in terms of localized phonons (Kittel 1996).

There is nothing particularly paradoxical about this: the crystal isn't 'really' a gas of phonons, it's just that certain states of the crystal have properties very similar to such a gas, and that treating these states as such is a great boon to analysis of crystal dynamics. This puts phonons and their ilk in good company, for a great many objects in science — such as animals, or rigid bodies — have to be understood in the same way. There are no perfectly rigid bodies, for instance (and they are certainly not part of the basic ontology of any fundamental physical theory), yet certain states of a many-particle system approximate the behaviour of 'ideal' rigid bodies extremely well, and so deserve the name. (See Wallace (2001a) for a more detailed discussion of this point.)

We shall adopt the same attitude to the particles of relativistic quantum

[2]See (Wilczek 1999) for an explicit statement of this consensus.

field theory: that is, we shall look for subspaces of the QFT Hilbert space in which the states have particulate properties. This will require us to formulate a definition of 'particle' and then to show that there are states of the QFT which approximately satisfy that definition; the rest of the paper is concerned with this task. First, though, we need to consider in which situations we would expect a QFT to appear particulate.

2.6 Particle regimes

The phenomenology of quantum field theory suggests two regimes in which we expect particle behaviour:

- **The non-relativistic limit**, in which the QFT appears to be described by slow-moving particles interacting by long-range forces;
- **The scattering limit**, in which particles begin widely separated, interact by short-range forces, and at late times are again found in widely separated states.

We shall be concerned almost exclusively with the second case, for reasons of mathematical tractability rather than on conceptual grounds: the analysis of relativistic fields via the methods of scattering theory is fairly well understood, whereas the process by which nonrelativistic quantum mechanics emerges as a limiting case of QFT is much more complicated. In the case of scattering theory, though, at times sufficiently long after (or before) the scattering event, the theory becomes very well approximated by a free quantum field theory. (This is intuitively plausible since for scattering theory to be applicable in the first place it is necessary that the nonlinear terms in the Hamiltonian constitute, after renormalisation, only a small perturbation to the free-field theory; for a much more careful discussion and justification, see Haag 1996.)

For this reason, our analysis henceforth will be restricted to free quantum theories (more specifically, to quantum theories of the form (1); this includes some sorts of background-field interactions).

3 Defining particles

In this section, we shall work out a definition of what properties a family of QFT states ought to have in order to count as 'particle' states. Since the idea of 'particle' is plainly at least connected to the concept of a localised state, we begin by considering how the latter states are to be defined in QFT.

3.1 Localised states in a field ontology

Which field-theory states are to count as localised?

In a QFT the idea of localisation must enter through the spatial localisation of the observables. The observables of the theory are defined via the field operators $\widehat{\phi}(\mathbf{x}, t)$ and $\widehat{\pi}(\mathbf{x}, t)$, so it is natural to define any given observable at time

t as being localised in a spatial region $\Sigma_i \subseteq \Sigma$ iff it is a function only of field operators of form $\widehat{\phi}(\mathbf{x}_i, t)$ and $\widehat{\pi}(\mathbf{x}_i, t)$ with all of the $\mathbf{x}_i$ in Σ_i.

But if defining localised observables is straightforward, defining localised states will prove decidedly less so. We might begin by trying:

> **Naive localisation**: A state $|\psi\rangle$ is localised in a spatial region Σ_i iff $\langle\phi|\,\widehat{O}\,|\phi\rangle = 0$ for any observable $\widehat{O}$ localised outside Σ_i.

This seems plausible when we compare it with the classical case: there a state is localised in Σ_i if $\pi(\mathbf{x}) = \phi(\mathbf{x}) = 0$ for any $\mathbf{x} \notin \Sigma_i$. But it is mathematically impossible for any state to satisfy it, for it implies that for any such $\mathbf{x}$, and for any $n \in \mathcal{Z}^+$,

$$\langle\psi|\,\widehat{\phi}^n(\mathbf{x}, t)\,|\psi\rangle = \langle\psi|\,\widehat{\pi}^n(\mathbf{x}, t)\,|\psi\rangle = 0. \tag{12}$$

But this would imply that $|\psi\rangle$ was a simultaneous eigenstate of $\widehat{\pi}(\mathbf{x})$ and $\widehat{\phi}(\mathbf{x})$, and these operators have no eigenstates in common. (The mathematics, bar some need to regularise to deal with operators defined at a point, is the same as for the nonrelativistic operators $\widehat{X}, \widehat{P}$, which are well-known to have no eigenstates in common.)

Physically it is easy to see what is happening here. The vacuum state of a field theory (which we will denote by $|\Omega\rangle$) is not 'nothingness', or 'empty space'; it is simply a slightly colourful way of describing the ground state of the field's Hamiltonian. In solid-state systems (which, recall, we are treating as field-theoretic systems like any other) this state is just the zero-temperature state of the solid, in which the atoms will not be at rest but will have zero-temperature fluctuations; the same will be true for the field excitations of a relativistic field theory.

This suggests, however, an alternative definition, first proposed (for spacetime regions $\mathcal{O}$, not spatial regions Σ_i) by Knight (1961):

> **Knight localisation:** a state is localised in a spacetime region $\mathcal{O}$ iff $\langle\phi|\,\widehat{A}\,|\phi\rangle - \langle\Omega|\,\widehat{A}\,|\Omega\rangle = 0$ for any observable $\widehat{A}$ localised outside the light cone of $\mathcal{O}$.

It is possible to find states satisfying this criterion (Knight 1961): take any unitary operator $\widehat{U}$ localised in $\mathcal{O}$, then the state $\widehat{U}\,|\Omega\rangle$ will be Knight-localised in $\mathcal{O}$.

However, Knight localisation differs in one important respect from the sort of localisation which we encounter in NRQM. In the latter, properties like 'is localised in $\mathcal{O}$' are treatable in the same way as properties like 'has energy E' or 'has momentum less than p': that is, we can define a projection operator whose intended interpretation is 'localised in $\mathcal{O}$', whose range is the space of all such states. This would be possible for Knight-localised states iff they form a subspace: that is, iff any superposition of two states Knight-localised in $\mathcal{O}$ is also Knight-localised in $\mathcal{O}$.

The fact that Knight-localised states do not have this property is a consequence of the Reeh-Schlieder theorem (Reeh and Schlieder 1961).

Reeh-Schlieder theorem: for any region $\mathcal{O}$, the set of all states generated by the action of operators localised within $\mathcal{O}$ upon the vacuum, spans the Hilbert space of the QFT.

(For a proof, and further discussion, see Haag 1996.) It follows[3] from the Reeh-Schlieder theorem that states Knight-localised at $\mathcal{O}$ span the entire state space, which rules out any possibility of a projector meaning 'localised with certainty in $\mathcal{O}$'.

It is easy to see — again by analogy to the solid state — why these problems occur. For in a generic solid-state system, atoms are coupled to their neighbours, and as a consequence the ground state of the system is highly entangled. This allows us (in principle) to exploit the long-range correlations between spatially separated subsystems of the field to produce any state by local operations within $\mathcal{O}$.

(To see this process in a far simpler system, consider the four- dimensional Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$, where $\mathcal{H}_A$ and $\mathcal{H}_B$ are each one-qubit (two-state) systems. The entangled states

$$|\phi\pm\rangle = \frac{1}{\sqrt{2}}(|1\rangle\otimes|1\rangle \pm |0\rangle\otimes|0\rangle) \tag{13}$$

are totally indistinguishable from one another when restricted to either subsystem (they both induce the reduced state $\rho = \frac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|)$ on each subsystem) but their sum $\frac{1}{\sqrt{2}}(|\phi+\rangle + |\phi-\rangle = |1\rangle\otimes|1\rangle$ is clearly distinguishable from both of them on either subsystem. Examples of this kind are analysed in rather more detail by Redhead (1995) and Clifton and Halvorson (2001).)

However, *in practice* the correlations due to vacuum entanglement usually drop off fast enough that using Knight-localised states to approximate states localised far from $\mathcal{O}$ requires prohibitively high-energy states. We can then use the following pragmatic criteria to characterise locality:

1. Effective localisation (qualitative form): A state $|\psi\rangle$ is *effectively localised* in a spatial region Σ_i iff for any function $\widehat{f}$ of field operators $\widehat{\phi}, \widehat{\pi}$, $\langle\psi|\,\widehat{f}\,|\psi\rangle - \langle\Omega|\,\widehat{f}\,|\Omega\rangle$ is negligibly small when $\widehat{f}$ is evaluated for field operators outside Σ_i, compared to its values when evaluated for field operators within Σ_i.

2. The effective localisation principle (ELP) (qualitative form: A subspace $\mathcal{H}$ of the QFT Hilbert space $\mathcal{H}_\Sigma$ obeys the ELP on scale L iff for any spatial region $\mathcal{S}$ large compared with L, a superposition of states effectively localised in $\mathcal{S}$ is effectively localised in effectively the same region.

[3] To see that it follows, we need only note that the unitary elements of a (bounded) operator algebra $\mathcal{A}(\mathcal{O})$ span $\mathcal{A}(\mathcal{O})$. This can be proved as follows: for any bounded Hermitian element $\widehat{H}$ of $\mathcal{A}(\mathcal{O})$, and any $t \neq 0$, $(it)^{-1}(\exp(it\widehat{H}) - \widehat{1})$ is a linear combination of unitary elements of $\mathcal{A}(\mathcal{O})$. As $t \to 0$, this sequence tends to $\widehat{H}$, hence $\widehat{H}$ is in the span of the unitary operators. To complete the proof, simply recall that any linear operator can be written as $\widehat{A} + i\widehat{B}$, where $\widehat{A}$ and $\widehat{B}$ are Hermitian.

These qualitative notions can be made precise in a number of ways, such as:

1. Effective Localisation (quantitative form): A state is L-*localised* in a region Σ_i, iff for any function $\widehat{f}$ of field operators $\widehat{\phi}, \widehat{\pi}$, $\langle\psi|\,\widehat{f}\,|\psi\rangle - \langle\Omega|\,\widehat{f}\,|\Omega\rangle$ falls off for large d like (or faster than) $\exp(-d/L)$, where d is the distance from Σ_i at which the function $\widehat{(f)}$ is evaluated. (Note that there is no difference, according to this definition, between a state L-localised at some spatial point $\mathbf{x}$ and a state L-localised in a region of size $\sim L$ around $\mathbf{x}$.)

2. ELP (quantitative form): A state obeys the ELP on scale L iff, for any 3-sphere $\mathcal{S}$ of radius $> L$, a superposition of states L-localised in $\mathcal{S}$ is L-localised in $\mathcal{S}$.

A subspace of states for which ELP holds on scale L can be treated — approximately — as possessing a well-defined concept of localisation and of "localised in Σ_i" projectors for regions large compared with L (these are constructed, for each such region Σ_i, by taking the projector onto the set of all states in $\mathcal{H}$ which are effectively localised in Σ_i; because of ELP, this set must be a linear space). Effectively, in such a subspace we are excluding enough states that for any sufficiently large Σ_i, we cannot construct states localised far from Σ_i using only those states localised within Σ_i.

It is still reasonable to ask: what good is effective locality? A state effectively localised in $\mathcal{A}$ can still in principle be distinguished from the vacuum via measurements made arbitrarily far away from $\mathcal{A}$. This question lies rather outside the scope of this paper (see Halvorson and Clifton (2001) and Wallace (2001b) for further discussion). Here we note only that such problems are by no means new to relativistic quantum theory. Even in non-relativistic quantum mechanics, there are in general no states which remain exactly localised in a finite region for any finite period of time — yet this does not seem to get in the way of the concept of localised particle in NRQM.

For the purpose of this paper, we shall treat effective localisation as 'good enough', and (since no particularly useful concept of exact localisation exists) will often drop the word 'effective', treating effectively localised states simply as localised.

3.2 What is a quantum particle?

Granted that a quantum field theory must be treated as being fundamentally about fields, what properties must a given state of a quantum field theory have in order to be deemed a particle state? It is instructive to start by considering the classical case: which classical field configurations (if any) could be described as particles? Here the answer seems obvious: the 'particle' configurations will be field configurations which are localised in a fairly small spatial region — localized blobs of field, in fact. Translated into quantum mechanics, this would make 'particles' just another name for the effectively localised states of the last section, provided that they were localised to sufficiently small regions.

However, this classical concept of particle is in one sense too weak to be appropriate for quantum theory. Classical wave-packets tend to spread out with time, becoming less localised — and hence, less 'particulate', whereas in non-relativistic quantum mechanics a state describing n particles at time t will continue to describe n particles at all other times — and even in relativistic quantum mechanics we wish to recover a notion of particle which is robust and time-independent provided the particles are far away from one another.[4]

Furthermore, the criterion that particles should be localised is in some sense also too strong for quantum mechanics. As the two-slit experiment reminds us, it is easy for a particle to enter a state which is nowhere near an eigenstate of position — in other words, nowhere near localised.

However, the two-slit experiment also suggests the correct quantum definition of particle. Although the experiment shows — by demonstrating interference of the particle wave — that a classical-particle picture isn't viable, it also shows that a classical-wave picture isn't viable either, because on measurement the particle is always found to be localised *somewhere*. To ensure within the formalism of quantum physics that this happens, it is enough to require the particle to be a linear superposition of states all of which are localised — then any measurement of particle position will always give a single answer. (I stress that this is intended to be an essentially interpretation-independent statement: I am not addressing the measurement problem here.)

These observations motivate our definition of a quantum particle:[5]

A space of one-particle states of size L (where L is small), written $\mathcal{H}_{1P}$, is a subspace of the QFT Hilbert space $\mathcal{H}_\Sigma$ such that

1. There is a basis for $\mathcal{H}_{1P}$, each member of which is a state L-localised at a point; equivalently, all states in $\mathcal{H}_{1P}$ are linear superpositions of such localised states.

2. $\mathcal{H}_{1P}$ satisfies the effective localisation principle on scale L.

3. $\mathcal{H}_{1P}$ is effectively preserved, on relevant timescales, by the dynamics of the field theory.

This definition is intentionally somewhat vague. The imprecision of the third criterion mirrors the way in which quasi-particles arise in solid-state physics — often the quasi-particles spontaneously decay, so that the one-particle subspace is not exactly preserved by the dynamics. However, provided that the decay time is long compared to other relevant timescales (such as the time taken by the quasi-particles to move between collisions) then the quasi-particles will

[4] We can find classical field theories which contain states like these — the solitons of the sine-Gordon equation are one example (Coleman 1985) — but in general they occur only in strongly non-linear theories, whereas here we are concerned with linear or nearly linear theories.

[5] It should be noted that this definition is closely related to the definition used in algebraic QFT, in which an n-particle state is defined as one which is able to trigger up to, but no more than, n detectors at a time. See Haag (1996, section II.4 and chapter VI) for more on this definition.

provide a useful concept with which to describe the field theory. As the decay time decreases there will come a point at which this concept ceases to be useful, but it would be a mistake to try to define this point exactly.

We have also made no attempt to be precise about the phrase 'where L is small': how small is small? In non-relativistic quantum mechanics, the answer is 'arbitrarily small': a (possibly overcomplete) basis can be constructed from states effectively localised in arbitrarily small regions of configuration space. (The set of all Gaussians of an arbitrary fixed width, for instance, will do nicely.) It will turn out, however, that this is not possible in quantum field theory: here there will turn out to be a minimum realizable size. It is reasonable to think of this as giving the 'size' of a particle: a particle's size is the size of the smallest region in which it can be localised.

Is it justifiable to be this vague in our definitions? A robust answer would be 'it works for quasi-particles, so why not?' More satisfactorily, we can recall that we are not looking for particles which can be added to the basic ontology of our theory (which, granted, does need precise definition); the basic ontology is and remains states of $\mathcal{H}_\Sigma$, or equivalently, wave-functionals on $\mathcal{S}$. Rather, we are just finding a good way to characterise certain states with interesting properties. Provided these states are picked out very accurately, there is no need to worry if the accuracy isn't perfect: we are simply looking for accurate, robust schemes by which we can approximate the dynamics of the theory and explain phenomena. (For a more extended, and somewhat more philosophical, defence of this use of approximate concepts in physics, see Wallace 2001a.)

In any case, it is the existence of an $\mathcal{H}_{1P}$ simultaneously satisfying (1), (2) and (3) which is in need of explanation. A space satisfying any given one of these clauses would not be particularly remarkable: for instance, given any collection of localised states we could construct a space satisfying (1) by taking their span, but then this space would not generally be preserved under time-evolution; or we could construct a space satisfying (3) by taking the collection of all states which are time-evolutes of our given collection, but then generally not all such states would be linear superpositions of members of the original collection. Furthermore, if our system satisfied (1) and (3) but not (2), we would have no guarantee that the concept of localisation would work for our particles as we need it to do in non-relativistic quantum mechanics and in scattering theory: specifically, we would have no guarantee of the existence of projections onto particles in a specific location.

In the next three sections, we will go about constructing states which fit the definition of a particle given above. Before embarking on this task, though, we should address an obvious objection: that we know perfectly well which states of a free QFT are the one-particle states, so all that is left to do is verify that the definition holds for these states.

The results of the ensuing calculation would, of course, confirm *that* free QFTs have one-particle sectors; however, it would not really answer the question of *why* they do. The more indirect approach used here is intended to give some insight into this second question.

4 Modal analysis of a free field

This section is a mathematical analysis of the structure of classical linear field theories; it is a common 'building block' for the two methods of reaching the one-particle subspace which will be developed in sections 5 and 6.

For the sake of mathematical rigour, this section makes some use of distribution theory (all such material can safely be skipped by any reader who does not get nervous upon sighting a Dirac delta function). The notation and terminology used is essentially that of Rudin (1991), especially chapters 6–7; in particular, use is made of Rudin's elegant 'multi-index' notation, in which

- an index α stands for an ordered n-tuple $(\alpha_1, \ldots, \alpha_n)$ with $\alpha_i \in \mathcal{Z}^+$;
- $D^\alpha := \left(\frac{\partial}{\partial x^1}\right)^{\alpha_1} \cdots \left(\frac{\partial}{\partial x^n}\right)^{\alpha_n}$;
- $|\alpha| := \alpha_1 + \cdots + \alpha_n$.

4.1 Required properties of $\mathcal{R}$

Recall that the free-field theories we are considering have the field equation (1), i. e.

$$\frac{\mathrm{d}^2}{\mathrm{d}t^2}\phi(\mathbf{x}, t) + \mathcal{R}\phi(\mathbf{x}, t) = 0.$$

We begin our analysis with a technical digression onto the operator $\mathcal{R}$ in this field equation. Specifically we will require the operator to have the following properties:

1. $\mathcal{R}$ is a continuous linear map from $C^\infty(\Sigma)$, the space of real smooth functions on Σ, to itself.[6]

2. $\mathcal{R}$ can be extended to a self-adjoint operator on (a dense subspace of) the space $\mathrm{L}^2(\Sigma)$ of square-integrable complex functions on Σ. (We shall identify $\mathcal{R}$ with its self-adjoint extension).

3. $\mathcal{R}$ is a local operator, in the sense that $\mathcal{R}f(\mathbf{x})$ depends only on the values of f in an arbitrarily small neighbourhood of $\mathbf{x}$.[7]

4. The spectrum of $\mathcal{R}$ is known to be real, since it is self-adjoint; we shall also require it to be positive and to be bounded below by a strictly positive eigenvalue. (In other words, zero is not an eigenvalue of $\mathcal{R}$; hence, $\mathcal{R}$ is invertible).

[6] 'Continuous' means 'continuous with respect to the topology on $C^\infty(\Sigma)$ induced by the family of semi-norms $p_N(f) = \sup\{|D^\alpha f(x)| : x \in \Sigma, |\alpha| \leq N\}$'; see Rudin (1991, pp. 34–36) for more on such topologies.

[7] Given the short-distance cutoff introduced in section 2.3 to make mathematical sense of interacting QFTs, the requirement of exact locality is not really necessary: it is enough to require that $\mathcal{R}f(\mathbf{x})$ depends significantly on the values of f only in a neighbourhood of width $\sim L_{\mathrm{cut}}$, where L_{cut} is the cutoff lengthscale; anticipating the later results of this section, this is to require that $\mathcal{R}$ is L_{cut}-local.

If the spectrum of $\mathcal{R}$ is discrete, $\mathcal{R}$ must have a complete set of eigenfunctions, orthonormal in the L^2 inner-product

$$\langle \phi, \psi \rangle \equiv \int_\Sigma \mathrm{d}^3\mathbf{x}\, \phi^*(\mathbf{x})\psi(\mathbf{x}); \tag{14}$$

we will denote a given such set as $f_k(\mathbf{x})$. (Note that since $\mathcal{R}$ is both real and self-adjoint we can always choose its eigenfunctions to be all real, though we shall not always do so.) The eigenvalue of f_k is denoted ω_k^2, with $\omega_k > 0$.

We will, in fact, take a somewhat schizophrenic attitude towards the discreteness (or otherwise) of the spectrum of $\mathcal{R}$: for conceptual analysis it will usually be convenient to take it as discrete, but in practical applications we will often want to take $\mathcal{R}$ to be a differential operator on R^3, in which case the spectrum is necessarily continuous. We shall therefore take the usual (if somewhat non-rigorous) physicist's step of assuming that moving from a discrete to a continuous spectrum is a purely technical matter involving no change in the conceptual situation.

Now, let $\mathbf{x}$ be any point on Σ; then we can define a linear functional $\mathcal{R}_\mathbf{x}$ on $C^\infty(\Sigma)$ by $\mathcal{R}_\mathbf{x} \cdot f = (\mathcal{R}f)(\mathbf{x})$; the continuity of $\mathcal{R}$ means that $\mathcal{R}_\mathbf{x}$ is continuous, hence is a distribution (generalised function) over $C^\infty(\Sigma)$. The following results are easy consequences of distribution theory and of the locality of $\mathcal{R}$:

1. Because $\mathcal{R}$ is local, each $\mathcal{R}_\mathbf{x}$ has support $\{\mathbf{x}\}$.

2. From theorem 6.25 of (Rudin 1991) we can deduce that (in a local chart at $\mathbf{x}$), we can find constants c_α and N such that $\mathcal{R}_\mathbf{x} = \sum_{|\alpha|\leq N} c_\alpha D^\alpha \delta_\mathbf{x}$, where $\delta_\mathbf{x}$ is a Dirac delta at $\mathbf{x}$.

3. From the continuity of $\mathcal{R}$, it follows that, in any local chart, we can find *functions* $c_\alpha(\mathbf{x})$ such that, for any $\mathbf{x}$ in the chart, $\mathcal{R}_\mathbf{x} = \sum_{|\alpha|\leq N} c_\alpha(\mathbf{x}) D^\alpha \delta_\mathbf{x}$.

4. From this, we deduce the (fairly obvious) fact that $\mathcal{R}$ is a differential operator.

If we follow the usual fiction of treating distributions as functions, we can (formally) define a function $\mathcal{R}(\mathbf{x}, \mathbf{y})$ by

$$\int_\mathcal{S} \mathrm{d}^3\mathbf{y}\, \mathcal{R}(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) \equiv \mathcal{R}_\mathbf{x} \cdot f; \tag{15}$$

hence

$$(\mathcal{R}f)(\mathbf{x}) = \int_\mathcal{S} \mathrm{d}^3\mathbf{y}\, \mathcal{R}(\mathbf{x}, \mathbf{y}) f(\mathbf{y}). \tag{16}$$

Again formally, we can think of this function as giving the matrix elements of $\mathcal{R}$ in a position basis, provided we remember that these elements are derivatives of delta functions.

It follows from the spectral theorem that

$$\mathcal{R}(\mathbf{x}, \mathbf{y}) = \sum_k \omega_k^2 f_k(\mathbf{x}) f_k^*(\mathbf{y}), \tag{17}$$

and that the kernels $\mathcal{R}^\lambda(\mathbf{x},\mathbf{y})$ of the operators $\mathcal{R}^\lambda$ are given by

$$\mathcal{R}^\lambda(\mathbf{x},\mathbf{y}) = \sum_k \omega_k^{2\lambda} f_k(\mathbf{x}) f_k^*(\mathbf{y}). \tag{18}$$

(Again, these kernels may well be delta-functions or other such distributions; they are not necessarily well-behaved functions.)

Fractional powers of $\mathcal{R}$ will become important later in the paper, and in general such operators will not be exactly local even if $\mathcal{R}$ is (the operator $\sqrt{m^2-\nabla^2}$, for instance, is known (Goodman and Segal 1965) to be anti-local, in the sense that for any function f, $\mathrm{supp}\sqrt{m^2-\nabla^2}f \cup \mathrm{supp} f$ is all of space except possibly for a set of points of measure zero) but they may be 'approximately local'. We define 'approximately local' as follows:

> An operator $\mathcal{R}$ is L-local iff its kernel $\mathcal{R}(\mathbf{x},\mathbf{y})$ drops off like $\exp(-|\mathbf{x}-\mathbf{y}|/L)$ as $|\mathbf{x}-\mathbf{y}|$ becomes large compared with L .

Informally, this means that while $\mathcal{R}f(\mathbf{x})$ does not just depend on the values of f in an arbitrarily small neighbourhood of $\mathbf{x}$, it does depend significantly on the values of f only in a neighbourhood of width $\sim L$. Note that there is a certain looseness in the definition (in the phrase 'large compared with'); purely mathematically, we could replace this with 'as $|\mathbf{x}-\mathbf{y}| \to \infty$' but clearly it would be against the spirit of the definition for (say) the kernel to start dropping off only once $|\mathbf{x}-\mathbf{y}| \gg 10^{30}L$.

In the next section we will prove approximate locality for an important subclass of $\mathcal{R}$ operators.

4.2 Euclidean-invariant $\mathcal{R}$

In this section we will consider an important sub-class of $\mathcal{R}$ operators: those which act upon R^3 and which are invariant under spatial translations and rotations. In this context we can establish the approximate locality of the $\mathcal{R}^\lambda$ operators.

The reason for requiring translation invariance is that we can work in Fourier space: any translation-invariant operator must have the exponential functions $\frac{1}{(2\pi)^{3/2}}\exp(i\mathbf{k}\cdot\mathbf{x})$ as its eigenfunctions and so from (17) we must have

$$\mathcal{R}(\mathbf{x},\mathbf{y}) = \frac{1}{(2\pi)^3}\int \mathrm{d}^3\mathbf{k}\, \exp(i\mathbf{k}\cdot(\mathbf{x}-\mathbf{y}))\omega^2(\mathbf{k}). \tag{19}$$

Since $\mathcal{R}(\mathbf{x},\mathbf{y})$ is a rotationally-invariant sum of derivatives of delta functions, it follows that the function $\omega^2(\mathbf{k})$ is a polynomial in $\mathbf{k}\cdot\mathbf{k}$. Formally, then, the integral (19) can be transformed to

$$\mathcal{R}(r) = \frac{1}{(2\pi)^2 ir}\int_{\mathrm{R}} \mathrm{d}k\, k\, \omega^{2\lambda}\exp(ikr) \tag{20}$$

where $r = |\mathbf{x}-\mathbf{y}|$ and where we have replaced $\mathcal{R}(\mathbf{x},\mathbf{y})$ with $\mathcal{R}(r)$ to indicate that $\mathcal{R}$ depends on $\mathbf{x}$ and $\mathbf{y}$ only through r. For positive λ at least, this integral

is divergent, indicating that $\mathcal{R}(r)$ is distributional; however, if the spectrum is unbounded then for sufficiently negative λ then the integral becomes convergent.

Now, general powers of polynomials are holomorphic except along branch cuts from the zeroes of the polynomial (for details, see any introductory complex analysis textbook; Priestley (1989) gives a clear exposition); furthermore, by assumption (4) of section 4.1 we know that the function $\omega^2(k)$ has no real zeroes. We are then able to evaluate (20) by contour integration, as follows (this is a simple generalisation of standard methods in QFT; see, e. g. , Peskin and Schroeder (1995) for a presentation of the method for $\mathcal{R} = m^2 - \nabla^2$): we construct branch points from each zero in the positive half-plane upwards parallel to the imaginary axis, and from each zero in the negative half-plane downwards parallel to the imaginary axis. We then deform the contour upwards towards $+i\infty$, wrapping it around each branch cut in the upper half-plane in the process. The integral is then a sum of integrals along both sides of each branch cut in the upper half-plane.

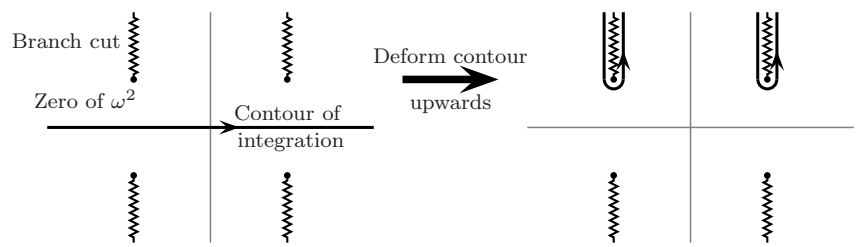

Now, suppose that there are N such branch cuts, starting at points (u_i, v_i) in the upper half-plane. Each integral along such a cut takes the form

$$I_i = \frac{\exp(-v_i r)}{(2\pi)^2 ir} \int_0^\infty \mathrm{d}\rho\, (u_i + i(v_i + \rho)) f_i(\rho) \exp(iu_i r) \exp(-\rho r), \tag{21}$$

where $f_i(\rho) = \lim_{\epsilon \to o}(\omega^{2\lambda}(u_i + \epsilon, v_i + \rho) - \omega^{2\lambda}(u_i - \epsilon, v_i + \rho))$. As $r \to \infty$ (or, more specifically, as r grows large compared to each $1/u_i$) the sum of all such terms will be dominated by I_{i_0}, where v_{i0} is the smallest of the v_i; in turn, for large r I_{i_0} will become proportional to $\exp(-v_{i_0} r)/r$.

This establishes that for sufficiently negative λ, $\mathcal{R}^\lambda$ is a $1/v_{i_0}$-local operator in the sense of section 4.1. Since *any* $\mathcal{R}^\lambda$ can be written as $\mathcal{R}^{\lambda-n}\mathcal{R}^n$ for arbitrarily large n, and since each $\mathcal{R}^n$ is strictly local, it follows that all powers of $\mathcal{R}$ are (at least) $1/v_{i_0}$-local.

In the case of Klein-Gordon theory, where $\mathcal{R} = m^2 - \nabla^2$, the only zero of the spectrum in the upper half plane is at $(0, m)$, hence powers of $\mathcal{R}$ are $1/m$-local in this case. $1/m$ is generally referred to as the *Compton wavelength*; we will extend this term to all translation-invariant $\mathcal{R}$, and define the Compton wavelength L_c of such an $\mathcal{R}$ as equal to v_{i_0}.

Given these results, in addition to axioms (1-4) of section 4.1, we will require $\mathcal{R}$ to satisfy one of the following:

5a. For all λ, $\mathcal{R}$ is L_c-local for some L_c; or

5b. $\mathcal{R}$ is rotationally and translationally invariant.

Of course, 5*b* implies 5*a*.

It might appear that solid-state systems do not satisfy 5b since the lattice structure violates translational and rotational invariance, but in fact the lattice only enters the observable results of the theory by imposing a short-distance cutoff, and hence (provided we work at lengthscales large compared with the cutoff) most solid-state systems may be treated as satisfying 5b.

4.3 Modes of the free field

Recall how to solve the free-field equation (1) by separation of variables: we try an ansatz of form $\psi(\mathbf{x}, t) = A(\mathbf{x})B(t)$; this gives

$$A(\mathbf{x})\ddot{B}(t) + B(t)\mathcal{R}A(\mathbf{x}) = 0. \tag{22}$$

Dividing through by $A(\mathbf{x})B(t)$ splits the equation into two terms, one independent of $\mathbf{x}$ and the other of t; this means that the equation can be solved only by finding solutions to the paired equations

$$\ddot{B} + \omega^2 B = 0; \tag{23}$$

$$\mathcal{R}A = \omega^2 A \tag{24}$$

where ω is to be determined. The second of these is simply the eigenfunction equation for $\mathcal{R}$. Each mode will have either exponential decay/growth (for $\omega^2 < 0$), or sinusoidal variation (for $\omega^2 > 0$), in time; our restriction to positive $\mathcal{R}$ eliminates the former case (this is the reason for this restriction) and we are left with a set of solutions of the form

$$\phi(\mathbf{x}, t) = f_k(\mathbf{x}) \cos(\omega_k t) \tag{25}$$

and

$$\phi(\mathbf{x}, t) = f_k(\mathbf{x}) \sin(\omega_k t). \tag{26}$$

(For the Klein-Gordon equation, the $f_\mathbf{k}$ are just proportional to sine and cosine functions $\sin(\mathbf{k} \cdot \mathbf{x}), \cos(\mathbf{k} \cdot \mathbf{x})$, with the possible values of $\mathbf{k}$ constrained by the boundary conditions and with $\omega_k^2 = m^2 + \mathbf{k} \cdot \mathbf{k}$.)

An arbitrary solution of the equations can be expressed as a sum of solutions of this form:

$$\phi(\mathbf{x}, t) = \sum_k \frac{1}{\sqrt{\omega_k}} \left(q_k f_k(\mathbf{x}) \cos(\omega_k t) + p_k f_k(\mathbf{x}) \sin(\omega_k t) \right), \tag{27}$$

so that a solution is given by the collection of real numbers $(q_k,\, p_k)$.

Since the space of solutions to the field equations is in one-to-one correspondence with the phase-space $\mathcal{P}$ (via $\phi(\mathbf{x}) \equiv \phi(\mathbf{x}, 0), \pi(\mathbf{x}) \equiv \dot{\phi}(\mathbf{x}, 0)$) we can regard (q_k, p_k) as coordinatizing $\mathcal{P}$: to be specific, we have

$$\phi(\mathbf{x}) = \sum_k \frac{1}{\sqrt{\omega_k}} q_k f_k(\mathbf{x}) \tag{28}$$

and

$$\pi(\mathbf{x}) = \sum_k \sqrt{\omega_k} p_k f_k(\mathbf{x}). \tag{29}$$

In fact, the choice of $\sqrt{\omega}_k$ factors in (27) means that they are canonical coordinates, in the sense that they obey the Poisson-bracket relations $\{q_k, q_{k'}\} = \{p_k, p_{k'}\} = 0; \{q_k, p_{k'}\} = \delta_{k,k'}$ (the proofs are straightforward and make use of the orthonormality of the f_k). In these coordinates the Hamiltonian (5) becomes

$$\mathrm{H} = \frac{1}{2} \sum_k \omega_k (p_k^2 + q_k^2). \tag{30}$$

Thus, subject to our restrictions on $\mathcal{R}$ at the start of section 4.1, any free-field theory can (as is of course well-known) be expressed as a sum of independent harmonic oscillators.

If we define $\alpha_k = \frac{1}{\sqrt{2}}(q_k + ip_k)$ we can rewrite (27) in the alternative form

$$\phi(\mathbf{x}, t) = \sum_k \frac{1}{\sqrt{2\omega_k}} \left(\alpha_k f_k(\mathbf{x}) \exp(i\omega_k t) + \alpha_k^* f_k^*(\mathbf{x}) \exp(-i\omega_k t)\right). \tag{31}$$

In fact, there is no real reason to restrict to real-valued eigenfunctions f_k: the expansion (31) is just as valid for complex eigenfunctions. Inverting it gives

$$\alpha_k = \frac{1}{\sqrt{2}} \int_{\mathcal{S}} \mathrm{d}^3\mathbf{x}\, f_k^*(\mathbf{x}) \left(\sqrt{\omega_k}\phi(\mathbf{x}) - i\frac{1}{\sqrt{\omega_k}}\pi(\mathbf{x})\right), \tag{32}$$

which (following the definition $\alpha_k = \frac{1}{\sqrt{2}}(q_k + ip_k)$ in the real-f_k case) suggests taking

$$q_k = \int_{\mathcal{S}} \mathrm{d}^3\mathbf{x} \left(\sqrt{\omega_k}\phi(\mathbf{x})\mathrm{Re} f_k(\mathbf{x}) + \frac{1}{\sqrt{\omega_k}}\pi(\mathbf{x})\mathrm{Im} f_k(\mathbf{x})\right) \tag{33}$$

and

$$p_k = \int_{\mathcal{S}} \mathrm{d}^3\mathbf{x} \left(-\sqrt{\omega_k}\phi(\mathbf{x})\mathrm{Im} f_k(\mathbf{x}) + \frac{1}{\sqrt{\omega_k}}\pi(\mathbf{x})\mathrm{Re} f_k(\mathbf{x})\right) \tag{34}$$

in the complex case. We can readily show that these are still canonical coordinates; the forms of (27–29) become slightly more complicated but H still has the form (30). This generalisation is useful in the Klein-Gordon equation, for instance, as it allows us to take $f_{\mathbf{k}} \propto \exp(i\mathbf{k}\cdot\mathbf{x})$, which is usually mathematically more convenient than working with sine and cosine functions.

The coordinate functions q_k and p_k have very simple time-dependence: from the Poisson brackets, we have $\dot{q_k} = \omega_k p_k$ and $\dot{p_k} = -\omega_k q_k$. Hence the time-evolution of the system is

$$q_k(t) = q_k(0)\cos(\omega_k t) + p_k(0)\sin(\omega_k t); \tag{35}$$

$$p_k(t) = p_k(0)\cos(\omega_k t) - q_k(0)\sin(\omega_k t), \tag{36}$$

or equivalently

$$\alpha_k(t) = \alpha_k(0)\exp(-i\omega_k t). \tag{37}$$

5 Particles through coherent states

The first method which we will use to construct particle states begins by constructing quantum-mechanical approximations to classically localised field states. In fact, it will turn out that these approximations are not particles, but they provide a natural stepping stone towards particles.

5.1 Harmonic oscillator coherent states

For a linear field theory, note that even in the classical case there are states which are localized to greater or lesser degrees. For instance, plane waves are (improper) classical states which are not localized at all, whereas we can construct fairly localized wave-packets. This is different from the case of particle quantum mechanics, where the classical states are perfectly localized and any loss of localization occurs only at the quantum level.

One way to construct a localized quantum state might be as follows: we begin by choosing a point in the classical phase space which is fairly localized (e. g. a fairly compact classical wave-packet) and then try to construct a quantum state which is concentrated around this point. Obviously, in the general case there is no unique way of approximating a phase-space point in quantum mechanics since precise localization in phase-space is not a well-defined quantum concept. In the case of a harmonic oscillator, however, there is a simple prescription(Glauber (1963); see Peres (1993) for a discussion) which generates approximations to phase-space points. If $\widehat{a}^\dagger$ is the creation operator for such an oscillator and $|0\rangle$ is its ground state, then the state

$$|\alpha\rangle = \exp(-|\alpha|^2/2)\exp(\alpha\widehat{a}^\dagger)\,|0\rangle \tag{38}$$

has the following properties:

1. It is a Gaussian in both position and momentum space;
2. It is centred around $q = \sqrt{2}\,\mathrm{Re}\,\alpha$ in position space, and $p = \sqrt{2}\,\mathrm{Im}\,\alpha$ in momentum space;
3. In both position and momentum space, the wave packet keeps its shape under time evolution (i. e. remains a Gaussian of the same width);
4. As time passes, the centres of the Gaussian in position and momentum space evolve as would the position and momentum of a classical harmonic oscillator with the same Hamiltonian.

Such a state is called a *coherent state*. Note that because of the one-to-one correspondence between phase space and the set of solutions to the dynamical equations, and because the coherent states track the classical evolution of phase points, we may equally well regard coherent states as quantum approximations to classical *solutions*.

Since the free field is (mathematically speaking) a collection of independent harmonic oscillators, these coherent states are an appropriate tool to construct quasi-classical states. The kth mode has creation operator $\widehat{a}_k^\dagger (= \frac{1}{\sqrt{2}}(q_k - ip_k))$, and hence a basis for the field Hilbert space $\mathcal{H}_\Sigma$ is given by the states created by successive actions of the different $\widehat{a}_k^\dagger$ on the vacuum.

A state localized around the kth mode would be

$$\widehat{D}_k(\alpha)\,|\Omega\rangle := \exp(-|\alpha|^2/2)\exp(\alpha\widehat{a}_k^\dagger)\,|\Omega\rangle \tag{39}$$

where the classical mode being approximated is $(\mathrm{Re}\,\alpha f_k/\sqrt{\omega_k}, \mathrm{Im}\,\alpha f_k\sqrt{\omega_k})$. Similarly, a classical state made from a superposition of modes may be quantum-mechanically approximated by successive application of $\widehat{D}_k$ operators to the vacuum, and the evolution of the classical state will be tracked by the quantum wave-packet.

It is vital to keep in mind the differences between classical and quantum concepts in what we are doing. Remember, we are constructing a quantum wave-packet, a complex functional on the space of field configurations, which is concentrated around a given point in configuration space. That point itself describes a classical wave-packet, that is, a real function on physical, three-dimensional space.

As time passes, the quantum wave-packet will not spread out across configuration space but will move around keeping its shape. Its centre will move through the configuration space according to the classical equations of motion, which will entail the spreading out through physical space of the *classical* wave-packet. Thus a coherent state will become less localized in physical space with time even though the quantum wave-functional keeps its shape (and, in particular, its width) in configuration space.

Now suppose the classical solution which we want to approximate is

$$\phi(\mathbf{x},t) = \sum_k \frac{1}{\sqrt{2\omega_k}}\left(\alpha_k f_k(\mathbf{x})\exp(i\omega_k t) + \alpha_k^* f_k^*(\mathbf{x})\exp(-i\omega_k t)\right), \tag{40}$$

and that it corresponds to the phase-space point (ϕ,π); then the corresponding coherent state is

$$|\mathcal{C}(\phi,\pi)\rangle = \prod_k \widehat{D}_k(\alpha_k)\,|\Omega\rangle \tag{41}$$

(the order in which the $\widehat{D}_k(\alpha_k)$ are applied is irrelevant as they all commute). Writing this out explicitly, we get

$$|\mathcal{C}(\phi,\pi)\rangle = \prod_k \left(\exp(-|\alpha_k|^2/2)\exp(-\alpha_k\widehat{a}_k^\dagger)\right)|\Omega\rangle \tag{42}$$

which may equally well be written as

$$|\mathcal{C}(\phi,\pi)\rangle = \exp\left(-\frac{1}{2}\sum_k |\alpha_k|^2\right)\exp\left(-\sum_k \alpha_k\widehat{a}_k^\dagger\right)|\Omega\rangle\,. \tag{43}$$

All we have used here is the elementary fact — applicable also to commuting operators — that a product of exponentials is equal to the exponential of the sum of their arguments.

Now if we take (ϕ, π) to be an element of $\mathcal{P}$ satisfying $\sum_k |\alpha_k|^2 = 1$, we can define

$$\widehat{a}^\dagger_{(\phi,\pi)} = \sum_k \alpha_k \widehat{a}^\dagger_k \tag{44}$$

and

$$\widehat{D}_{(\phi,\pi)}(z) = \exp(-|z|^2/2) \exp(z\widehat{a}^\dagger_{(\phi,\pi)}). \tag{45}$$

then we will have

$$|\mathcal{C}(\phi, \pi)\rangle = \widehat{D}_{(\phi,\pi)}(1)\, |\Omega\rangle\,. \tag{46}$$

Acting on the vacuum with $\widehat{D}_{(\phi,\pi)}(z)$ for higher values of $|z|$ creates coherent states localized around successively larger wave-packets; in this way, the action of $\widehat{D}_{(\phi,\pi)}(z)$ on the vacuum as we vary $|z|$ and hold (ϕ, π) fixed will map out a collection of states, whose span is a subspace of $\mathcal{H}_\Sigma$. It is easy to see that this space can also be spanned by those states generated from the vacuum by successive action of the $\widehat{a}^\dagger_{(\phi,\pi)}$ operator. Structurally there is a strong similarity to the subspaces created by the $\widehat{a}^\dagger_k$, although of course this new subspace is not preserved by time evolution.

5.3 Coherent states are not particles

Can the coherent states be regarded as quantum particles? Absolutely not. Although there are more and less localised non-relativistic quantum particles, and more and less localised coherent states, the two forms of localisation are wildly different. If (ϕ, π) and (ϕ', π') are phase-space points localised in different regions[8] of Σ then $(\phi + \phi', \pi + \pi')$ is a classical state which is non-localized in the sense that it is an extended field concentrated in two separated regions; if there are two spatially separated waves propagating on the surface of a pond then the excitations of that surface are in this sense nonlocal. And all coherent states are approximations to classical states: a coherent state formed around $(\phi+\phi', \pi+\pi')$ is non-localized only in the same sense as its classical progenitor.

If ψ and ψ' are localised wave-packets of a quantum particle, on the other hand, then $\psi + \psi'$ is nonlocalized in a wildly different way. Though we may be able to regard ψ and ψ' as approximately classical (approximating classical point particles), we cannot so regard $\psi + \psi'$. After all, if the nonlocal nature of these wave-packets could be understood in the classical way then the profound foundational problems of quantum nonlocality would never have arisen.

The coherent states offer us the possibility of constructing localised quantum states, but they are certainly not particles — localised or otherwise.

[8] Or even if just ϕ and ϕ', or just π and π', are localised in different regions.

5.4 Two linear structures

The argument why coherent states will not do as particles can be summed up by noting that the following diagram does not commute:

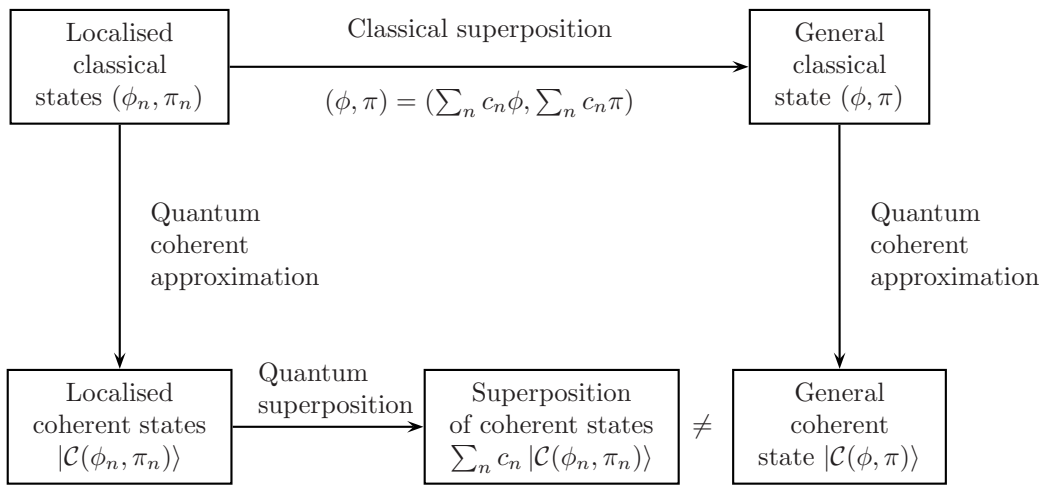

However, the strategy of approximating the classical states gives us a good method of constructing quantum states with given locality properties, and it would be useful to preserve that strategy in finding candidate particle states. Clearly, what is needed is a replacement of the 'quantum coherent approximation' arrows in the diagram with some other approximation, which still preserves localisation but which leads to a commutative diagram. Such a replacement would have to be a far weaker concept of approximation than that applicable to coherent states, preserving little more than localization properties. This is to be expected, in fact — non-local particle states are highly non-classical, so cannot be good approximations to any classical field state.

To find this replacement, consider the small-state limit of the coherent approximation: that is, consider coherent approximations to some state $(\lambda\phi, \lambda\pi)$, as $\lambda \to 0$. If (without loss of generality) we take the complex coordinates $\{\alpha_k\}$ of (ϕ, π) to satisfy $\sum_k |\alpha_k|^2 = 1$, then from (45), we find that such coherent approximations are given by

$$\begin{aligned} |\mathcal{C}(\lambda\phi, \lambda\pi)\rangle &= \exp(-\lambda^2/2) \exp(\lambda \widehat{a}^\dagger_{(\phi,\pi)}) \, |\Omega\rangle \\ &\simeq |\Omega\rangle + \lambda \, |\phi, \pi\rangle \end{aligned} \tag{47}$$

where we have defined

$$|\phi, \pi\rangle := \widehat{a}^\dagger_{(\phi,\pi)} \, |\Omega\rangle \, . \tag{48}$$

Now, if the classical state (ϕ, π) is non-zero (or just non-negligible) only in a region $\mathcal{A}$, then the state $|\phi, \pi\rangle$ which we have just defined is, plausibly, effectively

localised in $\mathcal{A}$: for it is a linear combination of $|\mathcal{C}(\phi,\pi)\rangle$ (which is a coherent approximation to (ϕ,π), and thus presumably localised in $\mathcal{A}$) and $|\Omega\rangle$ (which, trivially, is localised everywhere).

Of course, plausibility is at the moment all we have: as section 3.1 explained, effective localization is not in general a linear property. But in the following section we will calculate the actual localization properties of $|\phi,\pi\rangle$, and find that they are indeed localized in the same region as the classical state. For the moment note that if such states are localised correctly then they are precisely what we are looking for — for if we define $|k\rangle = \widehat{a}_k^\dagger |\omega\rangle := \left|f_k/\sqrt{2\omega_k}, 0\right\rangle$, we have (from 44) that

$$|\phi,\pi\rangle = \sum_k \alpha_k |k\rangle . \tag{49}$$

From this it follows immediately that the $|\phi,\pi\rangle$ states form a subspace, and that the linear structure on that subspace mirrors the linear structure of the classical states being approximated:

$$\widehat{a}^\dagger_{(A\phi+A'\phi',A\pi+A'\pi')} = A\widehat{a}^\dagger_{(\phi,\pi)} + A'\widehat{a}^\dagger_{(\phi',\pi')}. \tag{50}$$

5.5 Localization properties of $|\phi,\pi\rangle$ states

To investigate the localisation properties of the $|\phi,\pi\rangle$, we need to be able to calculate the expectation values of products of $\widehat{\phi}(\mathbf{x})$ and $\widehat{\pi}(\mathbf{x})$ operators, and the major tool we shall use will be knowledge of the (easily-calculated) commutators

$$\left[\widehat{\phi}(\mathbf{x}), \widehat{a}^\dagger_{\phi,\pi}\right] = \frac{1}{2}\left(\phi(\mathbf{x}) + i(\mathcal{R}^{-1/2}\pi)(\mathbf{x})\right) \tag{51}$$

and

$$\left[\widehat{\pi}(\mathbf{x}), \widehat{a}^\dagger_{\phi,\pi}\right] = \frac{1}{2}\left(\pi(\mathbf{x}) - i(\mathcal{R}^{1/2}\phi)(\mathbf{x})\right). \tag{52}$$

With these known, we can calculate expectation values by the usual method of moving the annihilation operators over to the right where they annihilate $|\Omega\rangle$. For instance:

$$\langle\phi,\pi|\,\widehat{\phi}(\mathbf{x})^2\,|\phi,\pi\rangle \equiv \langle\Omega|\,\widehat{a}_{\phi,\pi}\widehat{\phi}(\mathbf{x})^2\widehat{a}^\dagger_{\phi,\pi}\,|\Omega\rangle \tag{53}$$

$$= \langle\Omega|\,\widehat{\phi}(\mathbf{x})\widehat{\phi}(\mathbf{x})\,|\Omega\rangle + 2\left[\widehat{\phi}(\mathbf{x}), \widehat{a}^\dagger_{\phi,\pi}\right]\left[\widehat{\phi}(\mathbf{x}), \widehat{a}^\dagger_{\phi,\pi}\right]^* . \tag{54}$$

Subtracting off the vacuum expectation value (which is in general divergent, hence depends on whatever high-energy cutoff procedure we have chosen to use), we get

$$\langle\phi,\pi|\,\widehat{\phi}(\mathbf{x})^2\,|\phi,\pi\rangle - \langle\Omega|\,\widehat{\phi}(\mathbf{x})^2\,|\Omega\rangle = \phi(\mathbf{x})^2 + (\mathcal{R}^{-1/2}\pi)(\mathbf{x})^2. \tag{55}$$

In a similar way, we can calculate

$$\langle\phi,\pi|\,\widehat{\pi}(\mathbf{x})^2\,|\phi,\pi\rangle - \langle\Omega|\,\widehat{\pi}(\mathbf{x})^2\,|\Omega\rangle = \pi(\mathbf{x})^2 + (\mathcal{R}^{1/2}\phi)(\mathbf{x})^2, \tag{56}$$

$$\langle\phi,\pi|\,\frac{1}{2}\widehat{\pi}(\mathbf{x})^2+\frac{1}{2}\widehat{\phi}(\mathbf{x})\left(\mathcal{R}\widehat{\phi}\right)\mathbf{x})\,|\phi,\pi\rangle-\langle\Omega|\,\frac{1}{2}\widehat{\pi}(\mathbf{x})^2+\frac{1}{2}\widehat{\phi}(\mathbf{x})\left(\mathcal{R}\widehat{\phi}\right)\mathbf{x})\,|\Omega\rangle$$
$$=\frac{1}{2}\pi^2(\mathbf{x})+\frac{1}{2}\left(\mathcal{R}^{1/2}\phi\right)^2(\mathbf{x}), \tag{57}$$

etc. (The last expectation value is that of the energy density, i. e. the $(0,0)$ component of the stress-energy tensor.) In each case — and, it is easy to see, for any such expectation value — the expectation value is some function of the classical fields ϕ,π, modified by the action of some fractional power of $\mathcal{R}$. (In particular, the energy density is equal to the classical energy density up to the action of such operators.) Given that, at the end of section 4.2, $\mathcal{R}$ was required to satisfy either axiom 5b (which entails that all $\mathcal{R}^\lambda$ are L_c-local for some L_c) or axiom 5a (which requires this by fiat) it follows that

- If (ϕ,π) is localised in a region $\mathcal{A}$ then the difference of expectation values of $|\phi,\pi\rangle$ falls off like $\exp(-d/L_c)$ with distance d from $\mathcal{A}$.
- Hence, by the definition of effective localisation (in section 3.2), if (ϕ,π) is localised in a region $\mathcal{A}$ then $|\phi,\pi\rangle$ is effectively localised in the same region.
- In view of the linearity of the map $(\phi,\pi)\to|\phi,\pi\rangle$ which we have constructed between classical and quantum states, it follows that the space of all states $|\phi,\pi\rangle$ obeys ELP on scale L_c.

Note that (given the definition of L_c-localised states) any structure that the classical state has on scales smaller than L_c is likely to be disrupted by the action of $\mathcal{R}^\lambda$; in particular, for a classical state (ϕ,π) localised in a region small compared with L_c, the corresponding quantum state $|\phi,\pi\rangle$ will have little in common with the classical state other than being effectively indistinguishable from the vacuum at distances from the classical state which are large compared with L_c.

5.6 Particles at last

Let us review the process we have used to construct the $|\phi,\pi\rangle$ states. We have taken a classical wave-packet and constructed a coherent state around it. This state turns out to be expressible as the coherent state generated by a single creation operator, and in turn the action of that creation operator on the vacuum produces a state which is localised in the same region as the classical wave-packet (up to variations of size $\sim L_c$).

It is now easy to see that the following diagram commutes:

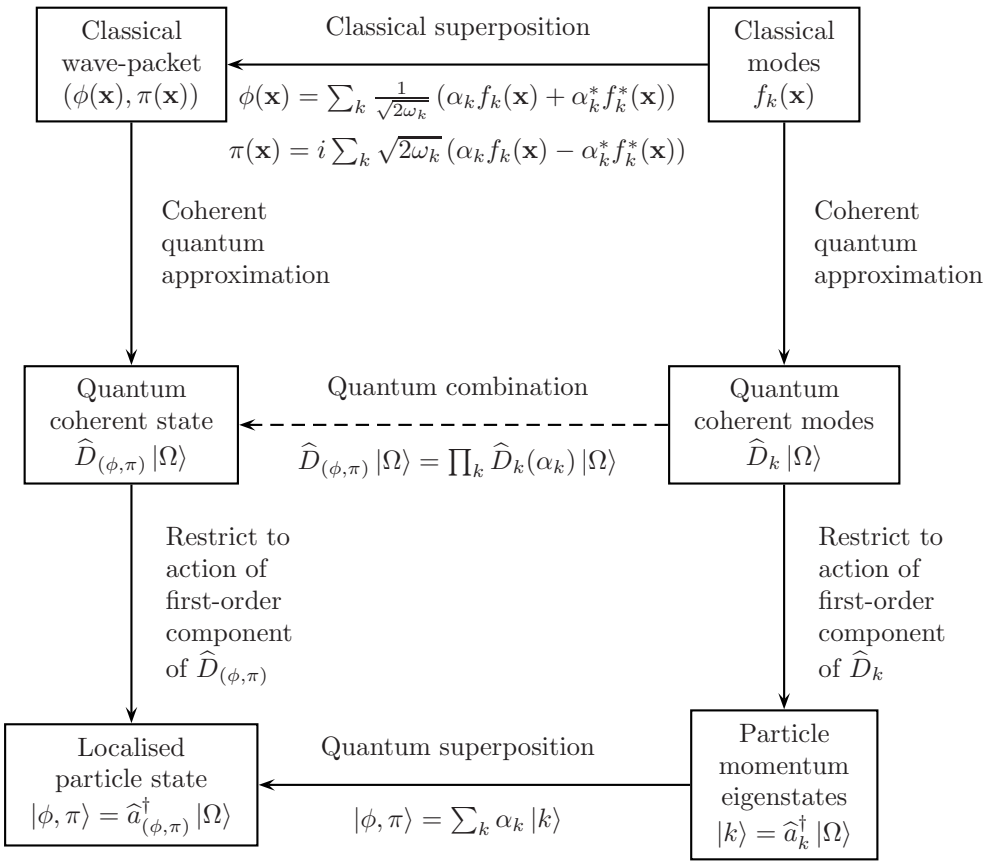

The important properties of the diagram are:

1. Moving *down* the diagram preserves L_c-localisation properties.

2. Moving from the first to the second row takes us from the classical to the quantum regime, but does not drastically change the nature of the states: states in the second row are good approximations of those in the first row, in the sense spelled out in sections 5.1–5.2 This is not true for the third row: the only sense in which $|\phi, \pi\rangle$ approximates (ϕ, π) is that they share the same localisation properties.

3. Moving leftward *across* the diagram corresponds to the combination of modes to make localised states. The middle (dashed) line is not a linear process, but the top and bottom lines both represent linear superposition. However, though mathematically very similar, these superpositions have physically very different meanings.

4. In the quantum superposition process, it is natural to consider complex weightings for the states being superposed. The (mathematically) equivalent process at the classical level provides a generalisation of the real-linear

superposition process for classical states, effectively equipping the classical solution space with a complex structure (of which more later).

With these results in hand, it is easy to verify that the space of $|\phi, \pi\rangle$ states is indeed a one-particle space in the sense of section 3.2. The localised states are constructed by beginning with classically localised wave-packets and moving down the diagram. The requirement that all states in the space are superpositions of localised ones follows from the equivalent property of classical phase space together with the commutativity of the diagram. The validity of the superposition principle among effectively localised states is a trivial consequence of the diagram's commutativity. And, crucially, the closure of the one-particle subspace under time-evolution follows once we observe that the $|k\rangle$ are energy eigenstates: this means that the projection operator onto the one-particle subspace commutes with the Hamiltonian, so the subspace must be preserved under time evolution. (Equivalently, closure under time evolution follows once we observe that the map

$$(\phi, \pi) \longrightarrow |\phi, \pi\rangle \tag{58}$$

commutes with time evolution.)

The essential property of the QFT which makes this whole process possible is its linearity: without the linearity, we would not have the *classical* linear structure whose interplay with the linear structure of $\mathcal{H}_\Sigma$ allowed our construction to proceed.

Henceforth, we will denote the space of all $|\phi, \pi\rangle$ by $\mathcal{H}_{1P}$.

6 Development of the particle concept

In this section we will analyse further the construction of particles presented above. We will examine the importance of the Compton wavelength, and develop the links between the linear structures on phase space and on Hilbert space; we will then use this analysis to give an alternative way of constructing the one-particle subspace.

6.1 Significance of the Compton wavelength

The results above imply that it is localisation on the scale of the Compton wavelength L_c, and not exact localisation, that is significant for particles. There is a straightforward physical reason for the significance of L_c: as mentioned in section 3.1, the vacuum state of any QFT is entangled (in the sense that field states in different spatial regions are entangled) and this entanglement cannot be removed from the non-vacuum states of the field without interfering with the field's structure at energy levels comparable to the cutoff energy (in other words, without going beyond the domain of validity of QFT). However, the correlations in the vacuum drop off with spatial distance, as can be seen from calculating quantities such as

$$\langle\Omega|\,\widehat{\phi}(\mathbf{x})\widehat{\phi}(\mathbf{y})\,|\Omega\rangle - \langle\Omega|\,\widehat{\phi}(\mathbf{x})\,|\Omega\rangle\,\langle\Omega|\,\widehat{\phi}(\mathbf{y})\,|\Omega\rangle$$
$$= \frac{1}{2}\mathcal{R}^{-1/2}\delta(\mathbf{x}-\mathbf{y}). \tag{59}$$

If $\mathcal{R}^{-1/2}$ is non-local on lengthscales of $\sim L_c$, then we can treat spatial regions separated by distances large compared with L_c as uncorrelated, but it makes rather little sense to talk about localisation on scales small compared with L_c.

This also gives us at least heuristic grounds to extend the concept of the Compton wavelength beyond Euclidean-invariant $\mathcal{R}$. As was shown in section 4.1, the locality of $\mathcal{R}$ requires it to have form

$$(\mathcal{R}f)(\mathbf{x}) = \sum_{|\alpha|\leq N} c_\alpha(\mathbf{x})D^\alpha f(\mathbf{x}), \tag{60}$$

with Euclidean-invariant $\mathcal{R}$ corresponding to each c_α being constant. Now, if we start with such a Euclidean-invariant $\mathcal{R}$ and introduce a very slow variation in its c_α (with 'very slow' meaning 'significant variation on lengthscales much longer than the Compton wavelength'), then we would expect the vacuum entanglement lengthscale to remain substantially unchanged, which in turn suggests that $\mathcal{R}^{-1/2}$ would remain non-local on the same lengthscales.

Of course, we are using physical intuition to conjecture results of a mathematical nature, and this conjecture is ultimately no substitute for rigorous results about the operators $\mathcal{R}^{-1/2}$ in the non-Euclidean-invariant case.

6.2 Field-particle duality

The map

$$(\phi,\pi) \longrightarrow |\phi,\pi\rangle \tag{61}$$

defines a vector-space isomorphism[9] between $\mathcal{H}_{1P}$ and the classical phase space $\mathcal{P}$; since it commutes with time evolution, it also defines an isomorphism between $\mathcal{H}_{1P}$ and the classical solution space. We can use this isomorphism to pull back the Hilbert-space structure (i. e., the complex structure and the inner product) from $\mathcal{H}_{1P}$ to phase space, and hence to solution space: thus, we gain a prescription by which we can regard $\mathcal{P}$ as a Hilbert space.[10]

It is instructive to give the precise forms for the complex structure (i. e. , the linear operator J corresponding to multiplication by i) and inner product $\langle\langle\,,\rangle\rangle$. on $\mathcal{P}$. We will express each in three ways:

[9] Strictly speaking the map takes phase space only into a proper subset of $\mathcal{H}_{1P}$, because the image of the map is not complete in the Hilbert-space norm on by the latter (induced from $\mathcal{H}_\Sigma$). Our relaxed attitude to this is due to our approach (in section 2.3) to renormalisation: we will regard $\mathcal{H}_{1P}$ as being cut off at some (very high) energy, thus making it finite-dimensional and removing the problem.

[10] The existence of a complex structure on the classical phase space has long been known, and in fact is a central part of Segal's (1964) approach to quantisation (briefly discussed in section 8.2).

1. In terms of q_k and p_k:

$$J(q_1, p_1; \ldots; q_k, p_k; \ldots) = (-p_1, q_1; \ldots; -p_k, q_k; \ldots);$$

$$\langle\langle\, (\phi, \pi), (\phi', \pi')\, \rangle\rangle = \frac{1}{2} \sum_k [(q_k q'_k + p_k p'_k) + i(q_k p'_k - p_k q'_k)]. \tag{62}$$

2. In terms of α_k:

$$J(\alpha_1, \ldots, \alpha_k, \ldots) = (i\alpha_1, \ldots, i\alpha_k, \ldots);$$

$$\langle\langle\, (\phi, \pi), (\phi', \pi')\, \rangle\rangle = \sum_k \alpha_k^* \alpha'_k. \tag{63}$$

3. Directly in terms of (ϕ, π):

$$J(\phi, \pi) = (-\mathcal{R}^{-1/2}\pi, \mathcal{R}^{1/2}\phi);$$

$$\begin{aligned}\langle\langle\, (\phi, \pi), (\phi', \pi')\, \rangle\rangle &= \frac{1}{2} \int_{\mathcal{S}} \mathrm{d}^3\mathbf{x}\, \Big(\phi(\mathbf{x})\mathcal{R}^{1/2}\phi'(\mathbf{x}) + \pi(\mathbf{x})\mathcal{R}^{-1/2}\pi'(\mathbf{x})\Big) \\ &\quad + \frac{i}{2} \int_{\mathcal{S}} \mathrm{d}^3\mathbf{x}\, \left(\pi(\mathbf{x})\phi'(\mathbf{x}) - \phi(\mathbf{x})\pi'(\mathbf{x})\right),\end{aligned} \tag{64}$$

or

$$\langle\langle\, (\phi, \pi), (\phi', \pi')\, \rangle\rangle = \langle \mathcal{R}^{1/4}\phi + i\mathcal{R}^{-1/4}\pi, \mathcal{R}^{1/4}\phi' + i\mathcal{R}^{-1/4}\pi' \rangle \tag{65}$$

where $\langle\, ,\, \rangle$ is the ordinary L^2 inner product on $\mathcal{S}$.

From (2) it is easy to confirm that J and $\langle\langle\, ,\rangle\rangle$ are preserved by time-evolution. The fractional powers of $\mathcal{R}$ which occur in (3) tell us that J and $\langle\langle\, ,\rangle\rangle$ are not strictly local.

Since $\mathcal{H}_{1P}$ and $\mathcal{P}$ are isomorphic, we can use this Hilbert-space description of $\mathcal{P}$ to provide a coordinatisation of $\mathcal{H}_{1P}$ itself. This gives us a sort of wave-function description, albeit one in which the complex structure and inner product are not locally defined. Such a description is an expression, in a sense, of wave-particle duality, with the same mathematical description applicable to the one-particle subspace of the quantum system, and to the whole classical system. Note, though, that the duality is critically dependent upon the linear structure of the solution space — hence we have no reason to regard field-particle duality as a general property of field theories, but only of linear (or nearly linear) ones.

The duality also implies that dynamics on phase space must be writeable in Schrödinger-equation form. Indeed, it is straightforward to check that Hamilton's equations

$$\dot{\phi} = \pi; \;\; \dot{\pi} = -\mathcal{R}\phi \tag{66}$$

are equivalent to

$$\frac{\mathrm{d}}{\mathrm{d}t}(\phi, \pi) = -J\mathcal{R}^{1/2}(\phi, \pi), \tag{67}$$

so that the Hamiltonian is the (mildly nonlocal) operator $\mathcal{R}^{1/2}$.

6.3 Alternative construction of particles: heuristic form

We have argued that it is the interplay between two Hilbert-space structures — the quantum-mechanical one and the one on the classical solution space — which leads to the emergence of particles, but it is perhaps somewhat obscure exactly how that interplay comes about. In this section and the next we will describe an alternative way to the particle subspace which possibly gives more insight into this question.

The basis of our new method is as follows: since the vacuum is significantly entangled only on lengthscales of order the Compton wavelength, we would expect that the action of a field observable like $\widehat{\phi}(\mathbf{x})$ on the vacuum would create a state differing from the vacuum only in the vicinity of $\mathbf{x}$ — that is, a state L_c-localised at $\mathbf{x}$. We might further expect that, if f and g are real functions which vanish outside some spatial region Σ_1, then the state

$$\int \mathrm{d}^3\mathbf{x} \left(f(\mathbf{x})\widehat{\phi}(\mathbf{x}) + g(\mathbf{x})\widehat{\pi}(\mathbf{x}) \right) |\Omega\rangle \tag{68}$$

would be L_c-localised in Σ_1. Furthermore, the space of all such states (68) is spanned by states of form $\widehat{\phi}(\mathbf{x})\,|\Omega\rangle$ and $\widehat{\pi}(\mathbf{x})\,|\Omega\rangle$ — that is, by states which we expect to be L_c-localised at a point.

It should be stressed that it is by no means obvious that these 'expectations' will be confirmed: the Reeh-Schlieder theorem reminds us that the link between the localisation properties of operators and of those states created by the action of those operators on the vacuum is rather subtle. Nonetheless, if they *are* confirmed then the space of states of the form (68) satisfies our first two criteria (on page 12) for a particle subspace: that the ELP holds for the space, and that the space is spanned by a basis of localised states. (And in fact they can be confirmed.)

No use has yet been made of the linearity of the classical solution space, so it is perhaps unsurprising that this property is essential to (amongst other things) ensure that our space is to satisfy the third criterion for a one-particle space: that the space is at least approximately preserved under time evolution. For the linearity of the classical field equations is equivalent to the requirement that the Hamiltonian is quadratic in the fields and their conjugate momentum, and this in turn entails that (writing $\widehat{U}(t) = \exp(-i\widehat{H}t)$ for the time-translation operator),

$$\widehat{U}(t)\widehat{\phi}(\mathbf{x})\widehat{U}^{\dagger}(t) \equiv \widehat{\phi}(\mathbf{x}) - it\left[\widehat{H}, \widehat{\phi}(\mathbf{x})\right] - t^2\left[\widehat{H}, \left[\widehat{H}, \widehat{\phi}(\mathbf{x})\right]\right] + \ldots \tag{69}$$

is a linear combination of $\widehat{\phi}$ and $\widehat{\pi}$ operators — hence, the time evolution of a state like $\widehat{\phi}(\mathbf{x})\,|\Omega\rangle$ is a state of form (68). This clearly entails the closure of the space of states of form (68) under time-evolution; hence, it is a one-particle space.

Let us develop the formal details of this sketch. We begin in classical mechanics: recall that a classical observable is a real functional on the phase space (so in one-particle mechanics the position observable assigns to each phase-space point its configuration-space position, etc.) It will be important to preserve, in the following, the distinction between the phase-space point (ϕ, π), which is a state of the classical field, and the classical observables $\phi(\mathbf{x})$ and $\pi(\mathbf{x})$, which are *functionals* on the space of field states: the action of $\phi(\mathbf{x})$ on a field state returns the field strength at the point $\mathbf{x}$.[11] To make this distinction clear, in this section we will distinguish classical observables by writing them with a bar on top of them: $\overline{\phi(\mathbf{x})}$, for instance. Thus, the observables $\overline{\phi(\mathbf{x})}$ and $\overline{\pi(\mathbf{x})}$ are defined by

$$\overline{\phi(\mathbf{x})}[(\phi, \pi)] = \phi(\mathbf{x}) \tag{70}$$

and

$$\overline{\pi(\mathbf{x})}[(\phi, \pi)] = \pi(\mathbf{x}). \tag{71}$$

The space of observables in any classical-mechanical system has a natural linear structure: $(\overline{A} + \overline{B})[\mathbf{v}] := \overline{A}[\mathbf{v}] + \overline{B}[\mathbf{v}]$. Furthermore, we can define time evolution for observables as follows: if, for any phase-space point $\mathbf{v}$, $\mathbf{v}(t)$ denotes where that phase-space point has moved to after time t, then we define

$$\overline{A}(t)[\mathbf{v}] := \overline{A}[\mathbf{v}(t)]. \tag{72}$$

If $\overline{H}$ is the classical Hamiltonian then we can write $\overline{A}(t)$ in the symbolic form

$$\overline{A}(t) = \exp\left(-t\left\{\overline{H}, \cdot\right\}\right)\overline{A}, \tag{73}$$

(which is to be understood as denoting the power-series expansion

$$\overline{A}(t) = \left(\sum_{n=0}^{\infty} \frac{(-t)^n}{n!}\left\{\overline{H}, \cdot\right\}\right)\overline{A} \tag{74}$$

where $\left\{\overline{H}, \cdot\right\}\overline{A} := \left\{\overline{H}, \overline{A}\right\}$.) This movement of the dynamics from the states to the observables is very similar, both conceptually and mathematically, to the move from the Schrödinger to the Heisenberg picture in quantum mechanics. It is discussed in more detail by Woodhouse (1991, p.20 *et seq.*).

Although we generally regard observables as real functionals, there is nothing to prevent us expanding the class of observables to include complex functionals, and we shall do so from here on. Any complex observable, of course, has the form $\overline{A} + i\overline{B}$, where $\overline{A}$ and $\overline{B}$ are real observables.

So far, everything we have said about observables applies to any classical system. If, however (as in the case of linear field theory) the classical phase space has a linear structure which is preserved under time evolution, then it

[11] Similarly, in particle mechanics there is a distinction between a particle's position $\mathbf{x}$ and the observables $\mathbf{x}^i$, which are functionals returning the ith component of a particle's position.

is possible to define *real-linear observables* as those real observables which are real-linear functionals; since the linear structure commutes with time evolution on phase space then it is easy to verify that it commutes with time evolution on the space of linear observables as well. In the specific case of classical fields, the linear observables are precisely those of form

$$\int_\Sigma \mathrm{d}^3\mathbf{x}\ \left(f(\mathbf{x})\overline{\phi}(\mathbf{x}) + g(\mathbf{x})\overline{\pi}(\mathbf{x})\right) \tag{75}$$

with f and g real functions. We will refer to any such observable as L_c-localised in a region Σ_1 if the functions f and g are L_c-localised in that region.

Since, as we have seen (in section 6.2), the classical phase space can be given the structure of a complex Hilbert space,[12] we can also define complex, conjugate-linear observables in the obvious way (i. e. all those complex observables which are real-linear and which satisfy $\overline{C}[J(\phi,\pi)] = -i\overline{C}[(\phi,\pi)]$.) It is easy to see that each such observable can be written as $\overline{C} = \overline{A} + i\overline{B}$, with $\overline{A}$ and $\overline{B}$ real-linear observables.

Now, it is a basic result of Hilbert space theory (the Riesz representation theorem) that the space of conjugate-linear functionals on a Hilbert space is isomorphic to that Hilbert space: in other words, that to each conjugate-linear functional Λ there corresponds a unique vector $\mathbf{u}$ such that, for any $\mathbf{v}$, $\Lambda[\mathbf{v}] = \langle \mathbf{u}, \mathbf{v}\rangle$. Hence, the space of complex conjugate-linear observables over the phase space of a classical linear field, conceived as a Hilbert space, is Hilbert-space-isomorphic to the phase space itself. This isomorphism takes the somewhat awkward form

$$(\phi,\pi) \longleftrightarrow \int_\Sigma \mathrm{d}^3\mathbf{x}\ \left((\mathcal{R}^{1/2}\phi(\mathbf{x}) - i\pi(\mathbf{x}))\overline{\phi}(\mathbf{x}) + (\mathcal{R}^{-1/2}\pi(\mathbf{x}) + i\phi(\mathbf{x}))\overline{\pi}(\mathbf{x})\right) \tag{76}$$

from which we can see that it preserves L_c-locality. It takes a far simpler form (albeit one in which its locality is obscured) if we use the α_k coordinates for (ϕ,π): if we define the observable $\overline{a}_k^*$ by $\overline{a}_k^*[\{\alpha_1,\ldots\alpha_n,\ldots\}] := \alpha_k$ then the isomorphism is

$$\{\alpha_1,\ldots\alpha_n,\ldots\} \longleftrightarrow \sum_k \alpha_k \overline{a}_k^*. \tag{77}$$

The move from classical to quantum mechanics can be thought of as an algorithm taking classical observables to quantum operators; it cannot be applied to all observables but in field theory it does apply to all of the real-linear observables, being generated by the maps

$$\overline{\phi}(\mathbf{x}) \longrightarrow \widehat{\phi}(\mathbf{x}) \tag{78}$$

and

$$\overline{\pi}(\mathbf{x}) \longrightarrow \widehat{\pi}(\mathbf{x}). \tag{79}$$

[12] As was mentioned briefly before (in footnote 9 on page 28) this is strictly correct only if we complete the classical phase space in the norm generated by its inner product.

The quantization map is linear, commutes with time evolution, preserves locality (trivially) and can be extended to complex conjugate-linear observables in the obvious way: $\overline{A} + i\overline{B} \longrightarrow \widehat{A} + i\widehat{B}$. As might be expected, it can be written as

$$\sum_k \alpha_k \overline{a}_k^* \longrightarrow \sum_k \alpha_k \widehat{a}_k^\dagger. \tag{80}$$

We have now moved from classical states, to classical observables, to quantum observables. Finally, we move to quantum states by applying the appropriate quantum observable to the vacuum state $|\Omega\rangle$. If the vacuum is entangled on lengthscales of L_c, then we expect (though it is not a priori guaranteed) that this map, too, is L_c-local.

This sequence of three linear isomorphisms, each commuting with time evolution, can be summarised in the following diagram.

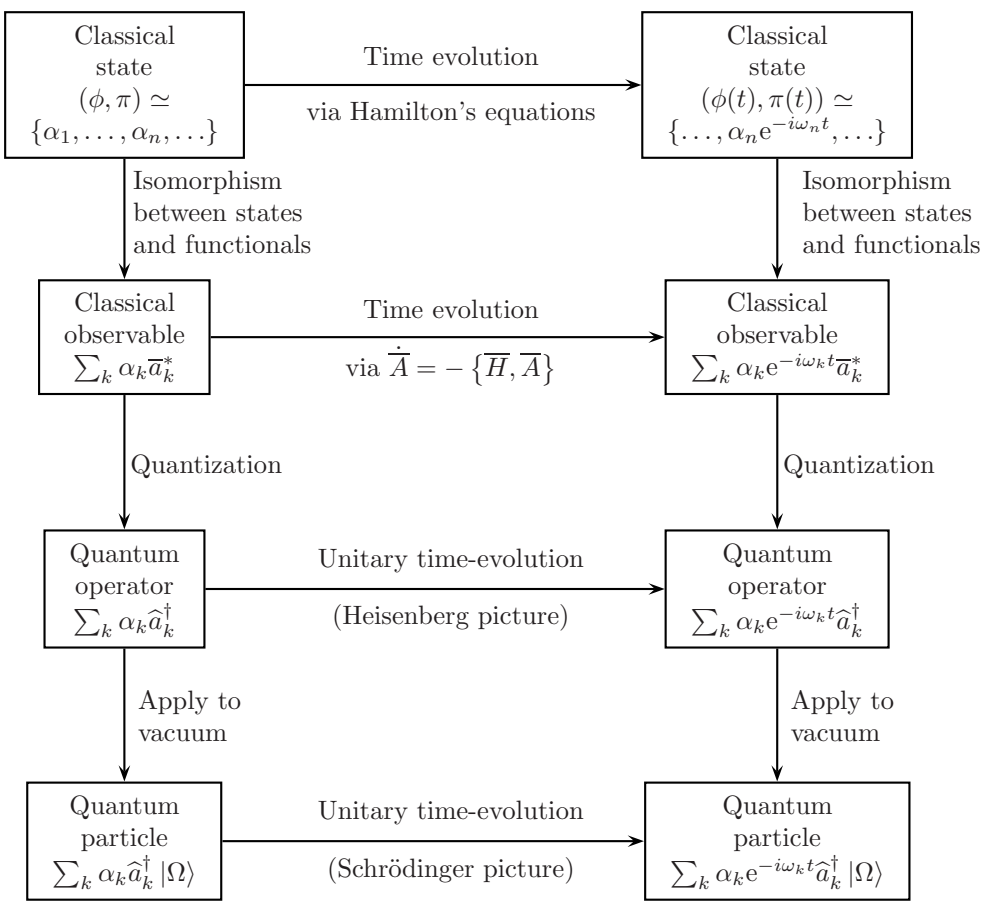

The important properties of the diagram are:

- Each map on the diagram is linear; hence, the linear structure on classical phase space transfers to the one-particle subspace.

- Nonetheless, the linear structures mean very different things: the linear structures on the classical and quantum observables are conceptually closely related to one another, but are conceptually different from the linear structure on the classical phase space, which in turn is conceptually different from the linear structure on the one-particle space.

- The vertical maps preserve L_c-localization; hence the one particle subspace satisfies the effective localisation principle on scale L_c and is spanned by states which are L_c-localised at a point. (This is true by definition between the first and second lines (cf. page 32), and trivial to show between the second and third lines. The substantive step is the one between the third and fourth lines — between quantum operators and quantum steps. It is made plausible by the observations of section 6.1 (that the vacuum is significantly entangled only on lengthscales of $\sim L_c$), and proved by the calculations of section 5.5.)

- The diagram commutes; in other words, time evolution commutes with each horizontal map. This ensures that states within the one-particle subspace remain in that space, i. e. that particle states are taken to particle states under time evolution.

Compared with the coherent-states method used earlier to construct particles, this new method has the advantage of showing more clearly how the linear structure transfers from classical to quantum states; however, it provides much weaker reasons for believing that the map between classical and quantum states preserves L_c-locality.

(Of course, even in the coherent-state method, the argument given for particle locality was only heuristic, relying on the assumption that if $|\mathcal{C}\rangle$ is a coherent state L_c-localised in some region, then $|\mathcal{C}\rangle - |\Omega\rangle$ is L_c-localised in the same region. Ultimately, the only way to check L_c-locality is by direct calculation.)

7 The Newton-Wigner representation of $\mathcal{H}_{1P}$

In this section, we will consider the so-called 'Newton-Wigner' definition of local states, which we will generalise from Lorentz-invariant QFTs to the linear QFTs discussed above.

7.1 Mathematical argument for the Newton-Wigner representation

Although the representation of $\mathcal{H}_{1P}$ given in the previous section is a sort of configuration-space representation, it is somewhat awkward to use compared to the configuration-space representations which we are accustomed to in non-relativistic quantum mechanics — essentially because of the non-local nature of the inner product and complex structure.

Purely for mathematical convenience, it would be useful to find a way of transforming our current representation into one where the inner product and complex structure are represented as local operations — such a transformation would inevitably have to be nonlocal itself, but since we realise that the particle states are unavoidably mildly nonlocal (on scales of L_c) this is not problematic. Finding such a transformation is straightforward: it is given by

$$(\phi, \pi) \to \mathcal{R}^{1/4}\phi - i\mathcal{R}^{-1/4}\pi, \tag{81}$$

or, in terms of the α_k coefficients, by

$$|\phi, \pi\rangle = \sum_k \alpha_k |k\rangle \to \sum_k \alpha_k f_k(\mathbf{x}). \tag{82}$$

It is straightforward to verify that in the new representation, the complex structure is just multiplication by i and the inner product is the usual L^2-product. Unless (ϕ, π) is localised in a region small compared with L_c then its localisation will not be significantly changed by the transformation; hence, the new representation is an equally accurate representation of the actual localisation properties of the quantum states. In the case of Klein-Gordon theory, this is known as the Newton-Wigner representation (Newton and Wigner 1949) and we will adopt this nomenclature for all such representations (i. e. , including QFTs other than Klein-Gordon theory).

If we denote the transformation from the phase-space representation of $\mathcal{H}_{1P}$ to the Newton-Wigner representation by $\mathcal{N}$, then its complex-linearity is equivalent to the statement

$$i\mathcal{N} = \mathcal{N}J. \tag{83}$$

Applying $\mathcal{N}$ to both sides of the Schrödinger equation (67), we get

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{N}(\phi, \pi) = -\mathcal{N}J\mathcal{R}^{1/2}(\phi, \pi) = -i\mathcal{R}^{1/2}\mathcal{N}(\phi, \pi) \tag{84}$$

so that the Hamiltonian in the Newton-Wigner representation is again $\mathcal{R}^{1/2}$. In the specific case of Klein-Gordon theory, we have $\mathcal{R}^{1/2} = (m^2 - \nabla^2)^{1/2}$, and in the non-relativistic limit this is approximately equal to $m - \nabla^2/2m$, so that we recover the non-relativistic free-particle Hamiltonian (up to an additive constant, the rest energy of the particle). It might appear that we have replaced a nonlocal operator with a local one, but for the nonrelativistic approximation to be valid we require that the wave-function contains only a negligible contribution from eigenstates which don't satisfy $|\mathbf{k}| \ll m$ — which in turn means that it must be localised in a region of size $\gg L_c$.

7.2 Conceptual significance of the Newton-Wigner representation

There is another way to motivate the Newton-Wigner representation, based more on conceptual than mathematical grounds. In Dirac's formulation of non-relativistic quantum mechanics, the configuration space wave-function is introduced by

$$\psi(\mathbf{x}) \equiv \langle \mathbf{x}|\psi\rangle\,, \tag{85}$$

where the states $|\mathbf{x}\rangle$ are the eigenstates of the position operator — so that the integral of the wave-function over a small region gives the amplitude for the particle to be in that region. This fits well with our description of a one-particle space in section 3.2 as a subspace spanned by well-localised states (indeed, in the non-relativistic case such states can be arbitrarily well localised) and it would be useful to have a similar representation of QFT particles.

To what extent does the Newton-Wigner representation provide this? The delta functions are certainly formally equivalent to position eigenstates, being perfectly localised in configuration space and forming an (improper) basis for the one-particle Hilbert space. But obviously they are not precisely localised in real space: if $|\mathbf{x}_{NW}\rangle$ is the abstract ket corresponding to a delta-function at $\mathbf{X}$, then we have

$$|\mathbf{x}_{NW}\rangle = \sum_k f_k^*(\mathbf{x})\,|k\rangle\,, \tag{86}$$

and it is easy to verify that (for instance)

$$\langle\mathbf{x}_{NW}|\,\widehat{\phi}(\mathbf{y})\widehat{\phi}(\mathbf{y})\,|\mathbf{x}_{NW}\rangle - \langle\Omega|\,\widehat{\phi}(\mathbf{y})\widehat{\phi}(\mathbf{y})\,|\Omega\rangle \tag{87}$$

is formally equal to

$$\frac{1}{2}\left[\mathcal{R}^{1/4}\delta(\mathbf{x}-\mathbf{y})\right]^2, \tag{88}$$

and hence is localised only within a region of size $\sim L_c$.

This is not a problem, however. Recall that even in non-relativistic quantum mechanics, no particle can actually be placed in an eigenstate of position (such states are not even in Hilbert space). Rather, such states are to be seen as an idealised mathematical limit of a sequence of successively better-localised states (see, e. g. , Cohen-Tannoudji, Diu, and Laloë (1977, pp. 100–105)). We can apply the same approach here, with the proviso that the successive terms of this sequence are only successively better localised in physical space (as opposed to in Newton-Wigner configuration space) up to a point: the point at which the terms are localised in regions which are not large in comparison with L_c. After that point the states will remain localised in a region of size $\sim L_c$ irrespective of how localised the Newton-Wigner wavefunction is. But providing that the wavefunctions which we are studying do not themselves vary on lengthscales comparable to L_c, we will only need those terms in the sequence which are large compared with L_c — hence, we may effectively use the concept of 'position eigenstate'.

Providing the Newton-Wigner wavefunction varies slowly on lengthscales of size L_c, we can regard the integral of the function over regions of this size as giving the amplitude to find it in a state effectively localised in the region. This then allows us to define an effective position operator:

$$\widehat{\mathbf{X}}_{NW} \equiv \int_\Sigma \mathrm{d}^3\mathbf{x}\,|\mathbf{x}_{NW}\rangle\,\langle\mathbf{x}_{NW}|\,. \tag{89}$$

This operator will perform the same task as the non-relativistic position operator, provided that we are not interested on scales small compared with the

minimum localisation scale: its expectation value will give the expected value of the particle's position upon measurement, and projections built from its spectrum correspond to projections onto the particle being in a given region (always assuming that region to be large compared with the minimum localisation scale L_c).

However, the relationship between this position operator and the Newton-Wigner position states is opposite from that which holds between non-relativistic position operators and position eigenstates (at least in some presentations). In the latter case, the position operator is our starting point and position eigenstates are position eigenstates *because* they are eigenstates of the position operator. In QFT, the Newton-Wigner position states are approximately local because of their relationship with the field observables, and the position operator is constructed from them.

7.3 "Pathological" features of the Newton-Wigner representation

The Newton-Wigner representation of Klein-Gordon particles is well known to have two apparently pathological features: eigenstates of position do not remain eigenstates under Lorentz transformations, and Newton-Wigner wavefunctions spread out superluminally. From the current perspective, both properties are curiosities rather than pathologies. The Newton- Wigner representation is not covariantly defined, so there is no mathematical reason to expect localisation to be totally unaffected by Lorentz boosts — but the boosts do not cause one-particle states which are effectively localised in a given region to stop being localised in that region. Similarly, Newton-Wigner wave-functions *do* spread faster than light — but the transformation between phase space $\mathcal{P}$ and the Newton-Wigner representation commutes with time evolution, so the Newton-Wigner state continues to represent a state effectively localised in the region occupied by a classical state which *is* propagating at subluminal speeds.

As such, the Newton-Wigner representation is a perfectly legitimate, and often very convenient, way of describing states in the one-particle subspace — but it doesn't give the exact truth of the matter as to where particles are localised, because there isn't one: particles are superpositions of field excitations with finite size, so any attempt to give a wave-function description down to arbitrarily small scales is inevitably going to be itself arbitrary at those scales.

Before ending this discussion, I should note that there is an alternative tradition (going back to (Segal 1964), and currently defended by Fleming (1996)) that takes Newton-Wigner-localised states as *by definition* localised. In this viewpoint, localisation of states is defined directly in terms of the Newton-Wigner position operator, rather than (as here) via the localisation properties of the operator algebra. Such an approach, by taking the Newton-Wigner position operator as having fundamental significance, must of course confront the apparent pathologies of the Newton-Wigner representation.

This paper, however, has as a starting assumption that localisation is defined via the operator algebra, and as such I will not discuss Fleming's approach. The reader interested in how (and if) the alternative tradition can come to terms

with the pathologies should see Fleming and Butterfield (1999) as well as the recent exchange of views between Fleming (2000) and Halvorson (2001).

8 Comparisons with other methods

This section is a brief comparison of the results of this paper with two other approaches to quantum field theory: the collision theory developed in the framework of algebraic QFT by Haag and Ruelle, and Segal's approach to quantizing linear systems.

8.1 Comparison with Haag-Ruelle collision theory

Haag-Ruelle theory (discussed in Haag (1996), and references therein) is an analysis of collision theory, and of the particle content of quantum fields, within the framework of algebraic QFT. Conceptually speaking this is essentially the same framework used here (although treated with significantly more mathematical care[13]): the field is taken as primary and particles are regarded as emergent concepts; thus, this paper's approach is complementary to, and not in conflict with, Haag-Ruelle theory.

A detailed description of the theory lies beyond the scope of this paper, but essentially it applies to any Poincaré-invariant QFT which has a discrete eigenvalue of the mass observable (other than zero) and for which the mass spectrum of the subspace orthogonal to the vacuum has a lower bound away from zero (so that it implies, for instance to massive Klein-Gordon theory). And it establishes that at asymptotically early and late times,

1. The QFT behaves as a free, massive QFT;
2. Any state which lies in the N-particle sector of that free theory can set off at most N widely separated detectors.

Note that the theory applies to interacting as well as to free theories; however, it only applies at asymptotically early and late times, and implies that any QFT is effectively free at such times; thus, it entails a particle interpretation only for effectively free fields (which is the case described above). Its virtue over the approach of this paper is its generality (it applies to fermionic as well as bosonic fields, and to composite as well as 'elementary' particles) and its ability to infer the particle properties of the field directly from its mass spectrum, without recourse to concepts of field-particle duality. However, for this very reason the close relationship between classical solutions and one-particle states is somewhat obscure in the Haag-Ruelle analysis.

At first sight, the generality of the Haag-Ruelle analysis might seem to undermine the argument of section 6.2 that the particle concept is essentially bound up with the linearity of the field theory in question. However, the analysis itself tells us that the existence of a massive asymptotic limit of a field theory, and

[13] See Wallace (2001b) for a discussion of the relationship between algebraic QFT and the prima facie less rigorous QFT described here and used in mainstream physics.

the existence of a discrete nonzero mass, are really equivalent statements. In the context of the present paper, the Haag-Ruelle analysis should remind us that the linear theories we are studying should be regarded as effective, asymptotic limits of interacting field theories, and not as fundamental fields in their own right: for instance (though technically it lies rather beyond the scope of this paper since it deals with fermionic particles) we would expect QCD to be asymptotically analysable in terms of 'proton' and 'neutron' fields, but these bear little resemblance to the fundamental — and strongly interacting — quark fields of QCD.

(Technically, the analysis of this paper might seem more general than the Haag-Ruelle theory since the latter is confined to Poincaré-invariant QFTs. However, we can extend Haag-Ruelle theory to general translation-invariant QFTs by generalising its spectral condition to require that the 4-momentum spectrum includes a discrete surface in 4-momentum space which nowhere intersects the origin (we get this for free by Lorentz invariance in a relativistic theory, once we have specified a discrete mass eigenstate). It is at best unclear how to apply Haag-Ruelle theory to QFTs without translational symmetry; it should, though, be acknowledged that the above analysis is also not entirely satisfactory in such situations as we have only a heuristic argument for the approximate locality of non-translation-invariant $\mathcal{R}^\lambda$ operators — recall the discussion of section 6.1.)

8.2 Comparison with Segal quantization

It is instructive to compare the field-quantization approach given above with the quantization program developed by Segal (1964) and others (see Saunders (1992) for a foundational discussion). In Segal's approach, we *begin* by choosing a complex structure J on the phase space of a linear field,[14] which we require to have certain properties (specifically, for any states u, v, we require J to satisfy $\Omega(Ju, Jv) = \Omega(u, v)$, and $\Omega(Ju, u) \geq 0$, where $\Omega(\cdot, \cdot)$ is the symplectic form on $\mathcal{P}$). We then use J to define an inner product on $\mathcal{P}$, given by

$$\langle\langle\, u, v \,\rangle\rangle \equiv \Omega(Ju, v) - i\Omega(u, v). \tag{90}$$

This (following norm-completion) makes $\mathcal{P}$ into a Hilbert space (the 'one-particle Hilbert space'), on which we can easily show that the classical dynamics are unitary. The Fock space generated from that Hilbert space is then taken to be the Hilbert space of the QFT, and there is a very natural definition of the action of the field operators upon that space.

This is effectively the reverse procedure to ours. In Segal's approach, we begin with the Hilbert-space structure on phase-space, declare it to be the one-

[14] There is another approach very closely related to Segal quantization, in which instead of choosing a complex structure, we complexify the phase space and then choose a subspace of (formally) half the dimension of the complexified phase space. This process is at least formally equivalent to choosing a complex structure, though there are some infinite-dimensional technicalities to take care of: see Wald (1994) for an exposition of this method, and especially p. 46 of his book for the relationship with the Segal method.

particle Hilbert space, and then build the full QFT Hilbert space $\mathcal{H}_\Sigma$ from that structure; in ours, we begin with $\mathcal{H}_\Sigma$, identify a subspace which deserves to be called the 'one-particle' space, and then use it to give a Hilbert-space structure to $\mathcal{P}$.

This highlights an important conceptual difference between Segal quantization and the 'canonical' method of quantization used here. In our approach, the method used to quantize a (bosonic) field theory is essentially the same irrespective of whether that theory is linear: throw the theory into Hamiltonian form and look for a map from field observables to operators on some Hilbert space, such that Poisson brackets become commutators. Then if the classical theory is linear, analysis of the quantum theory reveals the one-particle subspace. In Segal's approach, on the other hand, the linearity is fundamental to the entire quantization algorithm, and it is most unclear how nonlinear field theories are to be accommodated.

The attitude which one may take towards this difference depends on one's confidence in the mathematical and conceptual status of nonlinear QFTs — in particular, the status of the infinities which arise in the dynamics of such theories, and of the renormalisation process used to tame them. If (with the author; see Wallace 2001b) one takes a relaxed attitude to these infinities, then the Segal process becomes an elegant curiosity: a nice way to understand linear quantization, but of little relevance to general QFTs. If, on the other hand, one regards the infinities as a pathological problem in QFT, then (as Saunders (1992) has argued) it may be a virtue of the Segal approach that it forces us to look for radical ways of reformulating the quantization of nonlinear field theories.

It is also possible to see Segal quantization as pointing to a different view of the relationship between field and particle than the one defended here. As mentioned above, in Segal quantization it is necessary to begin by choosing a complex structure on the classical phase space, and different choices of structure lead to one-particle spaces with different properties. If one believes that there is a 'right' choice of complex structure, then there is also a 'right' one-particle space, and that space is a fundamental building block when assembling the field Hilbert space; this is at least suggestive of a more even relationship between field and particle concepts than the field-is-primary viewpoint advocated above.

On the other hand, it is also reasonable to regard the choice of a complex structure as just being a matter of taste, in which case the importance of the one-particle subspace is downgraded. One argument in favour of this attitude is that in general it is very unclear what rule could pick out the 'right' complex structure: there is a unique prescription in the case of a time-independent linear QFT (this includes the QFTs considered above) but in, say, the presence of a time-dependent potential, or in a curved spacetime, then there does not appear to be any general rule by which one complex structure could be preferred over another. This is, of course, a variant of the last of the arguments which (in section 2.4) led us to regard field as primary in the first place.

9 Conclusion

We have shown that the one-particle subspace of a free, bosonic, massive quantum field emerges very naturally as a consequence of that field's classical linearity, and that the interplay between classical and quantum linear structures is such as to give those particle states the right sort of locality properties.

It is important to remember that this sort of construction of particle spaces is not intended to be any sort of explanation as to why the world around us is observed to be particulate. At best it is able to show how such observations are compatible with the validity of QFT — the question of why the actual world appears to be in a particulate state is much subtler. It presumably requires consideration of decoherence theory (see Anglin and Zurek (1996) for some ideas along these lines), and more generally it involves the measurement problem.

It is also important to keep in mind the limitations on our construction. Although these results should extend straightforwardly to multi-component fields, they do not apply to the important cases of fermionic, and massless, fields. In the former case we would expect the particle concept to be more, not less, important, due to the Pauli exclusion principle; but the fact that spacelike-separated Fermi field operators anticommute rather than commute suggests that we can no longer visualise the QFT Hilbert space in terms of wave-functionals on configuration space. In the case of massless fields, we expect the particle concept to be much more subtle:[15] QED, for instance, possesses not only regimes in which the electromagnetic field is analyzable in terms of free photons, but regimes in which it should be thought of as a classical field, and others in which it simply provides a long-range force between nonrelativistic particles.

This brings up perhaps the most substantial limitation of this paper's approach to particles — and indeed of any approach which analyses particles in terms of asymptotic behaviour. Although it may well be the case that there is no particulate description of high-energy phenomena except at asymptotically early and late times, atomic and solid-state physics — not to mention biology and chemistry — provide us with a wealth of situations in which 'particles' (specifically, electrons and atomic nuclei) are strongly interacting and yet still maintain their own particle character. Analysing this situation would require a deep understanding of the nonrelativistic limit of massive and massless QFTs, and lies far beyond this paper's scope.

Nonetheless, despite the limitations of this approach it is satisfying to find that field-particle duality can be understood in the context of a field ontology, and intriguing to observe the elegant and somewhat subtle ways in which the various linear and complex structures present intertwine to make this understanding possible.

Acknowledgements— I would like to thank Simon Saunders, whose own work on the foundations of QFT prompted my interest in the subject; and Jeremy Butterfield for

[15] This is supported by the situation in algebraic QFT, in which development of a satisfactory collision theory for fields without a mass gap is at a much more rudimentary stage; see Haag (1996) and Buchholz (2000) for recent discussions.

many useful comments on earlier drafts of this paper. I have also benefitted greatly from conversations with Ian Aitchison, James Binney, Katherine Brading, Harvey Brown, Keith Burnett, Peter Morgan, and Andrew Steane.

References

Anglin, J. and W. H. Zurek (1996). Decoherence of quantum fields: Pointer states and predictability. *Physical Review D 53*(12), 7327–7335.

Buchholz, D. (2000). Algebraic quantum field theory: a status report. Plenary talk given at the XIIIth International Conference on Mathematical Physics, London. Available online at http://xxx.arxiv.org/math-ph/0011044.

Clifton, R. and H. Halvorson (2001). Entanglement and open systems in algebraic quantum field theory. *Studies in the History and Philosophy of Modern Physics 32*, 1–31. Available online at http://xxx.arXiv.org/abs/quant-ph/0001107.

Cohen-Tannoudji, C., B. Diu, and F. Laloë (1977). *Quantum Mechanics*. New York: Wiley.

Coleman, S. (1985). *Aspects of Symmetry: selected Erice lectures*. Cambridge: Cambridge University Press.

Fleming, G. (1996). Just how radical is hyperplane dependence? In R. Clifton (Ed.), *Perspectives on Quantum Reality*, pp. 11–28. Dordrecht: Kluwer Academic Publishers.

Fleming, G. (2000). Reeh-Schlieder meets Newton-Wigner. *Philosophy of Science 67 (Proceedings)*, S495–S515.

Fleming, G. and J. N. Butterfield (1999). Strange positions. In J. N. Butterfield and C. Pagonis (Eds.), *From Physics to Philosophy*, Chapter 6, pp. 108–165. Cambridge: Cambridge University Press.

Glauber, R. J. (1963). Coherent and incoherent states of the radiation field. *Physical Review 131*, 2766.

Goodman, R. and I. E. Segal (1965). Anti-locality of certain Lorentz-invariant operators. *Journal of Mathematics and Mechanics 14*, 629–638.

Haag, R. (1996). *Local Quantum Theory: Fields, Particles, Algebras*. Berlin: Springer-Verlag.

Halvorson, H. (2001). Reeh-Schlieder defeats Newton-Wigner: on alternative localisation schemes in relativistic quantum field theory. *Philosophy of Science 68*, 111–133.

Halvorson, H. and R. Clifton (2001). No place for particles in relativistic quantum theories? Available online at http://xxx.arXiv.org/abs/quant-ph/0103041.

Kittel, C. (1996). *Introduction to Solid State Physics*. New York: Wiley.

Kittel, C. and C.-Y. Fong (1987). *Quantum Theory of Solids*. New York: Wiley.

Knight, J. M. (1961). Strict localisation in quantum field theory. *Journal of Mathematical Physics 2*, 459.

Marsden, J. E. and T. S. Ratiu (1994). *Introduction to Mechanics and Symmetry*. New York: Springer-Verlag.

Newton, T. D. and E. P. Wigner (1949). Localized states for elementary particle physics. *Reviews of Modern Physics 21*, 400.

Peres, A. (1993). *Quantum Theory: Concepts and Methods*. Dordrecht: Kluwer Academic Publishers.

Peskin, M. E. and D. V. Schroeder (1995). *An introduction to Quantum Field Theory*. Reading, Massachusetts: Addison-Wesley.

Priestley, H. A. (1989). *Introduction to Complex Analysis*. Oxford: Oxford University Press.

Redhead, M. (1995). More ado about nothing. *Foundations of Physics 25*, 123.

Reeh, H. and S. Schlieder (1961). Bemerkungen zur unitäräquivalenz von lorentzinvarianten feldern. *Nuovo Cimento 22*, 1051.

Rudin, W. (1991). *Functional Analysis*. New York: McGraw-Hill.

Saunders, S. (1992). Locality, complex numbers and relativistic quantum theory. *Philosophy of Science Association 1992 1*, 365–380.

Segal, I. (1964). Quantum fields and analysis in the solution manifolds of differential equations. In W. Martin and I. Segal (Eds.), *Analysis in Function Space*, Cambridge, Massachusets. MIT Press.

Wald, R. M. (1994). *Quantum Mechanics on Curved Space-Time and Black Hole Thermodynamics*. Chicago: University of Chicago Press.

Wallace, D. (2001a). Everett and structure. To appear in *Studies in the History and Philosophy of Modern Physics*; available online at http://xxx.arXiv.org/abs/quant-ph/0107144 or from http://philsci-archive.pitt.edu.

Wallace, D. (2001b). In defence of naiveté: the conceptual status of Lagrangian quantum field theory. Available online at http://www.arxiv.org/abs/quant-ph/0112148 or from http://philsci-archive.pitt.edu.

Weinberg, S. (1995). *The Quantum Theory of Fields*, Volume 1. Cambridge: Cambridge University Press.

Wilczek, F. (1999). Quantum field theory. *Reviews of Modern Physics 71*, S85–S95. Available online at http://www.arXiv.org/abs/hep-th/9803075.

Woodhouse, N. M. J. (1991). *Geometric Quantization*. Oxford: Oxford Science Publications.